

JCS90 U.S. PTO
09/27/00

09.29.00

A

PTO/SB/05 (2/98) (modified)

Approved for use through 9/30/2000, OMB 0651-0032
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE

**NEW UTILITY
PATENT APPLICATION
TRANSMITTAL**

*(only for new nonprovisional applications under
37 CFR 1.53(b))*


Attorney Docket Number **5382**

First Named Inventor **William R. Caid**

Total Pages in this Submission **143**

Express Mail Label No. **EL550795024US**

U.S. PTO
09/27/00

APPLICATION ELEMENTS		ACCOMPANYING APPLICATION PARTS	
<p>1. <input type="checkbox"/> Fee Transmittal Form (in duplicate) <input type="checkbox"/> Check Enclosed</p> <p>2. <input checked="" type="checkbox"/> Specification <i>(preferred arrangement set forth below)</i> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> Descriptive Title of the Invention <input checked="" type="checkbox"/> Cross Reference(s) to Related Case(s) <input checked="" type="checkbox"/> Statement Regarding Fed sponsored R & D <input checked="" type="checkbox"/> Background of the Invention <input checked="" type="checkbox"/> Brief Summary of the Invention <input checked="" type="checkbox"/> Brief Description of the Drawing(s) <input checked="" type="checkbox"/> Detailed Description <input checked="" type="checkbox"/> Claim or Claims <input checked="" type="checkbox"/> Abstract of the Disclosure </p> <p>3. <input checked="" type="checkbox"/> Drawing(s) <i>(when necessary per 35 USC 113)</i></p> <p>4. Oath or Declaration <ul style="list-style-type: none"> a. <input type="checkbox"/> New Declaration <input checked="" type="checkbox"/> Copy of declaration from parent application b. <input type="checkbox"/> Copy from a prior application (37 CFR 1.63(d)) <i>(for continuation/divisional with Box 17 completed)</i> <ul style="list-style-type: none"> i. <input type="checkbox"/> DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 CFR 1.63(d)(2) and 1.33(b). </p> <p>5. <input checked="" type="checkbox"/> Incorporation by Reference <i>(useable if Box 4b is checked)</i>. The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby incorporated by reference therein.</p>		<p>6. <input type="checkbox"/> Assignment & Assignment Recordation Cover Sheet</p> <p>7. <input type="checkbox"/> Certified Copy of Priority Document(s) <i>(if foreign priority is claimed)</i></p> <p>8. <input type="checkbox"/> Information Disclosure Statement & PTO-1449 <input type="checkbox"/> Copies of IDS Citation(s)</p> <p>9. <input type="checkbox"/> Preliminary Amendment</p> <p>10. Small Entity Statement <ul style="list-style-type: none"> <input type="checkbox"/> New Statement enclosed <input type="checkbox"/> Statement filed in prior application. Status still proper and desired </p> <p>11. <input checked="" type="checkbox"/> Return Postcard</p> <p>12. <input checked="" type="checkbox"/> Communication to Examiner</p> <p>13. <input type="checkbox"/> _____</p> <p>14. <input type="checkbox"/> _____</p> <p>15. <input type="checkbox"/> _____</p> <p>16. <input type="checkbox"/> _____</p>	
ADDRESS TO: Assistant Commissioner for Patents Box Patent Application Washington, D.C. 20231			
17. If a CONTINUING APPLICATION, check appropriate box and supply the requisite information below and in a preliminary amendment:			
<input checked="" type="checkbox"/> Continuation <input type="checkbox"/> Divisional <input type="checkbox"/> Continuation-in-part (CIP) of prior application No: <u>08/931,927</u>			
Prior application information: Examiner: <u>J. Rhodes</u> Group/Art Unit: <u>2762</u>			
18. CORRESPONDENCE ADDRESS			
NAME	Robert R. Sachs Fenwick & West LLP		
ADDRESS	Two Palo Alto Square		
CITY	Palo Alto	STATE	CA
COUNTRY	U.S.A.	TELEPHONE	(415) 875-2410
Name (Print/Type)	Robert R. Sachs		Registration No. (Attorney/Agent) 42,120
Signature			Date Sept. 27, 2000

A9000/00000/SF/5019226.1

CONTEXT VECTOR GENERATION AND RETRIEVAL

Inventors

William R. Caid, Pu Qing, Joel L. Carleton, and David J. Sudbeck

Prepared by:

Robert R. Sachs

Registration No. 42,120

Fenwick & West LLP

Two Palo Alto Square

Palo Alto, CA 94306

Express Mail No.: EL5507950254US

Filed: September 27, 2000

**CONTEXT VECTOR
GENERATION AND RETRIEVAL**

INVENTORS

William R. Caid, Pu Qing, Joel L. Carleton, and David J. Sudbeck

5

CROSS-REFERENCE TO RELATED APPLICATION

This application is a continuation of U.S. Serial No. 08/931,927, filed September 17, 1997, now U.S. Pat. No. x,xxx,xxx, which was a continuation of U.S. Serial No. 08/322,313 filed October 13, 1994, which was a continuation-in-part of U.S. application Serial No. 08/124,098 filed by Caid et al., on September 20, 1993 (hereinafter, Caid et al.), which was 10 abandoned in favor of file wrapper continuation Serial No. 08/561,167 now U.S. Pat. No. 5,619,709 for "System and Method of Context Vector Generation and Retrieval." All of these related applications are incorporated by reference in their entirety.

BACKGROUND OF THE INVENTION

1. Field of the Invention

15 This invention relates to vector-based meaning-sensitive information storage and retrieval systems, and more particularly to an improved system and method for generating and retrieving context vectors that represent high-dimensional abstractions of information content. This invention also relates to storage, indexing, and retrieval of image data, and more particularly to a method and system for generating and retrieving context vectors that represent high-dimensional 20 abstractions of information in images.

2. Description of the Related Art

Conventional methods of record storage and retrieval generally involve storage of all records word for word and then searching for key words in the records using inverted indexes.

The key word searches are performed by doing a complete search through all of the contents of the data base that contain a list of query words. Such systems have no knowledge, for example, that "car" and "automobile" represent nearly the same meaning, so the user must include this information by using a complex and difficult-to-formulate query. Some systems try to solve this

5 problem by a built-in thesaurus, but such systems lack "meaning sensitivity" and miss many subtleties of meaning association, such as the fact that "car" is closer to "road" than to "hippopotamus".

There is currently much research and development in the field of neural networks. A neural network consists of a collection of cells and connections among cells, where every

10 connection has an associated positive or negative number, called a weight or component value. Each cell employs a common rule to compute a unique output, which is then passed along connections to other cells. The particular connections and component values determine the behavior of the network when some specified "input" cells are initialized to a set of values. The component values play roughly the same role in determining neural network behavior as a

15 program does in determining the behavior of a computer.

Prior art for document retrieval includes systems using variable length lists of terms as a representation, but without meaning sensitivity between terms. In such systems, pairs of terms are either synonyms or not synonyms.

So-called "vector space methods" can capture meaning sensitivity, but they require that

20 the closeness of every pair of terms be known. For a typical full-scale system with over 100,000 terms, this would require about 5 billion relationships—an impractical amount of information to obtain and store.

Methods have also been proposed for searching with fixed-length vectors. However,

such methods require work on the order of at least the square of the sum of the number of

25 documents and the number of terms. This is impractical for a large corpus of documents or terms.

A document retrieval model based on neural networks and capturing some meaning sensitivity has been proposed. However, a search in such models requires multiplications for twice the product of the number of document and the number of keywords for each of a plurality of cycles.

- 5 Koll in "WEIRD: An Approach to Concept-Based Information Retrieval," SIGIR Forum, vol. 13, no. 4, Spring 1979, pp. 32-50, discloses a retrieval method using vector representations in Euclidean space. The kernel or core used by Koll are non-overlapping documents. This results in rather small dimensional vectors on the order of seven values. Vectors are generated from the core documents based on whether or not a term appears in a document. As an
10 alternative, Koll suggests starting with a kernel of terms which never co-occur.

In addition, analysis of image subject content is a time-consuming and costly operation. This analysis is often required for the identification of images of interest in existing image data bases and the routing and dissemination of images of interest in a real-time environment. The conventional approach is to rely upon human intellectual effort to analyze the content of images.

- 15 It would be desirable to reliably translate image data into representations that would enable a computer to assess the relative proximity of meaning among images in a database.

It would be desirable to have a computing system that can derive accurate, efficient, and manageable representations of images for later recall, retrieval, and association.

SUMMARY OF THE INVENTION

- 20 The present invention is directed to a system and method for generating context vectors for use in a document storage and retrieval system. A context vector is a fixed-length series of component values representative of meaning or content. Geometric relationships among context vectors are representative of conceptual relationships among their associated items. Thus, two information items having similar meaning or content have similarly-oriented context vectors,

while items having dissimilar meaning or content have orthogonal context vectors. Similarity between items may be measured by calculating the dot product of the associated context vectors.

Context vectors may be associated with words, terms, documents, document portions, queries, images, quantitative data, people, or any other type of information item. This use of context vectors provides a context-sensitive information retrieval, routing, and visualization system based on learned similarity of meaning.

The present invention provides a context vector generation scheme that uses a neural network operating on a training corpus of records. Resulting vectors are relationship-based, formed by the proximity of words to one another in the training records. Relationships among context vectors are deterministic, so that a context vector set has one logical solution, although it may have a plurality of physical solutions. No human knowledge, thesaurus, synonym list, knowledge base, or conceptual hierarchy, is required.

Record storage according to the present invention is performed by inputting each record, which may be any type of information item, in machine-readable form into a processing system. If a record is textual in nature, uninteresting words may be removed from consideration prior to further processing. A learning law is then applied to each word (or analogous component) of the record, which assigns context vectors in accordance with word proximity. The learning law employs a technique of “windowed co-occurrence” wherein a fixed-size moving window is applied throughout the document, and words within the window (neighbor words) may exert “influence” on neighbor words in accordance with mutual co-importance. Such “influence” is constrained to avoid convergence, or collapse, of context vectors. Once context vectors are established for all components of a record, they are combined to form a summary vector for the record. This summary vector is then normalized. It represents the overall meaning or content of the record.

Summary vectors of records may be clustered to reduce searching time, by forming a tree of clustered nodes. A centroid is computed for each node based on the average of the normalized summary vectors within a defined group. Thus, records with similar information content occupy

the same cluster. Information content of the cluster is represented by the centroid vector. The node tree may be traversed to provide a speedy method of locating relevant records in the data base.

Once the context vectors are determined, records may be retrieved using a query interface
5 that allows a user to specify content terms, Boolean terms, and/or document feedback. The basic searching method involves converting the user's query into a context vector (a query vector). Queries may also be based on example documents, selected paragraphs, sentences, or words. The query vector is then combined with the normalized summary vectors (or with node centroids) to locate the records having the closest vectors. Retrieved records may be displayed in
10 order of vector proximity, which corresponds to relative relevance to the query. Rank ordering by proximity prevents information overload to the user (unlike conventional Boolean search methods, where search results may include, for example, 500 documents, without any indication of which documents are likely to be the most relevant). In addition, the most relevant portions of retrieved records may be highlighted if desired.

15 The system may also employ relevance feedback, whereby the user specifies which of the retrieved documents are most helpful. A new search may then be performed using the summary vector for the specified documents. This technique reduces the time required for searches and improves system effectiveness.

Furthermore, the techniques of the present invention facilitate visualization of textual
20 information by translating context vectors into visual and graphical representations. Thus, a user can explore visual representations of meaning, and can apply human visual pattern recognition skills to document searches.

Finally, the present invention provides a language-independent information representation scheme. Thus, it may be used to perform English queries on foreign text for retrieval (and vice
25 versa) without the need for prior translation or interpretation.

The present invention is also directed to a method and system for generating context vectors associated with images in an image storage and retrieval database system. In this aspect

of the present invention, context vectors are associated with information elements, or features, derived by performing wavelet transformations at a plurality of points on each electronically stored image in the database. The transformations provide orientation-sensitive spatial frequencies on the images at a variety of orientation/ frequency combinations. These features are
5 combined to form image feature vectors or “image vocabulary” elements analogous to words in text.

A prototypical subset of feature vectors, or atoms (also called information elements), are derived from the set of feature vectors to form an “atomic vocabulary.” In one embodiment, the prototypical feature vectors are derived by using a vector quantization method (e.g., self
10 organization) in which a vector quantization network is also generated.

The atomic vocabulary is used to define images in the database or any new image in electronic computer-readable form. As above, a wavelet transformation is performed at a plurality of sample points on the image to generate feature vectors representing the image. The generated feature vectors are mapped to the closest atoms in the atomic vocabulary using the
15 vector quantization network. Thus, new images are defined in terms of the established atomic vocabulary.

In one embodiment, a “stop list” of high-frequency, low-information, feature vectors is also generated. The stop list can be used to remove high-frequency, low-information, feature vectors when using the atomic vocabulary to represent images.

20 In order to quantify conceptual relationships among atoms in the atomic vocabulary (and the images they variously represent), context vectors are employed. A context vector is associated with each atom in the atomic vocabulary. A learning law is applied to modify the context vectors as a function of the proximity of the atom to other atoms in the image and the frequency of occurrence of the atom in the image database.

25 Once the context vectors are established, the context vectors associated with the atoms that define an image are combined to form a summary vector for the image. The summary vector represents the overall meaning or content of the image.

In one embodiment, summary vectors of images are stored in clusters to reduce searching time. Images with similar information content occupy the same cluster. In one embodiment, textual index terms are associated with images in the database, and are automatically assigned to new images. Thus, textual queries can be used to retrieve images.

- 5 Images are retrieved using any of a number of query methods (e.g., images, image portions, vocabulary atoms, index terms). The query is converted into a query context vector. A dot product calculation is performed between the query vector and the summary vectors to locate the images having the closest vectors. Retrieved images are displayed in order of vector proximity, which corresponds to relative relevance to the query. In one embodiment, retrieved
10 images are broken into sub-portions and the most relevant portions matching the query vector are highlighted in the image.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1A is a block diagram of the hardware elements of the present invention.

FIG. 1B is a block diagram of the training system of the present invention.

FIG.s 2A through 2F show an example of window definition.

FIG. 3 is a flowchart of the training process of the present invention.

FIG. 4 is a block diagram of the process of determining a summary vector for a document.

FIG. 5 shows sample icon attributes.

20 FIG. 6 is a sample display for document visualization.

FIG. 7 shows examples of icons representing associations.

FIG. 8 shows an example of user navigation through vector space.

FIG. 9A is a sample display for document visualization.

FIG. 9B is a sample display showing icon qualification.

25 FIG. 9C is a sample display showing an open icon.

- FIG. 10 is a flowchart of the retrieval method of the present invention.
- FIG. 11 is a flowchart of the method of document visualization.
- FIG. 12 is a chart of ranked eigenvalues derived from a set of approximately 15,000 280-dimensional context vectors.
- 5 FIG. 13 is a chart of ranked eigenvalues derived from a set of approximately 15,000 512-dimensional context vectors.
- FIG. 14 is graphical representation of the hierarchical a data object definition and organization in a Docuverse system.
- 10 FIG. 15 is a graphical representation of possible visual degrees of freedom for icons used in the Docuverse system.
- FIG. 16 is a graphical representation of the coordinate selection and information display used in the Docuverse system.
- 15 FIG. 17 is a graphical representation of icon displays used in the Docuverse system.
- FIG. 18 is a graphical representation of using a “point of view” in the Docuverse system.
- FIG. 19 is an example of specifying coordinates in the Docuverse system.
- FIG. 20 is an example of qualifying and selecting data in the Docuverse system.
- FIG. 21 is an example of changing the point of view in the Docuverse system.
- FIG. 22 is an example of selecting icon words and attributes in the Docuverse System.
- 20 FIG. 23 is an example of zooming in the Docuverse system.
- FIG. 24 is an example of inspecting data in the Docuverse system.
- FIG. 25 is a data flow diagram of context vector generation in the MatchPlus system.
- FIG. 26 is a continuation of FIG. 25.
- FIG. 27 is a data flow diagram of routing query generation in the MatchPlus system.
- FIG. 28 is a data flow diagram of retrieval in the MatchPlus system.
- 25 FIG. 29 is a simplified block diagram of a computer hardware and software implementation of an image retrieval system in accordance with the present invention .

FIG. 30 is a high level flowchart of an image retrieval system in accordance with the present invention.

FIG. 31 is a flowchart of one embodiment of the method for generating an atomic vocabulary.

5 FIG. 32 is a flowchart of one embodiment of the method for generating a "stop list."

FIG. 33 is a flowchart of one embodiment of the method for representing a new image using the atomic vocabulary.

FIG. 34 is a flowchart of the method for initializing context vectors.

FIG. 35 is a flowchart of one embodiment of the boot-strapping procedure.

10 FIG. 36 is a pictorial representation of FIG. 35.

FIG. 37 is a flowchart of one embodiment of the method for generating a summary vector for an image.

FIG. 38 is a high level flowchart of the clustering method of the present invention.

FIG. 39 is a flowchart of one embodiment of the clustering method.

15 FIG. 40 is a high-level flowchart of one embodiment of the indexing method.

FIG. 41 is a flowchart of one embodiment of the retrieval method.

FIG. 42 is a flowchart of one embodiment of the highlighting method.

FIG. 43 is a pictorial representation of highlighting as used in the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Section I. Context Generation and Retrieval

Referring now to FIG. 1A, there is shown a block diagram of a typical implementation of a system 100 in accordance with the present invention. The user supplies queries to system 100 via input device 111. Central processing unit (CPU) 107 runs software program instructions, stored in program storage 112, which direct CPU 107 to perform the various functions of system 100. In the embodiment illustrated herein, the software program is written in the C programming language and runs under the UNIX operating system. Each of these languages may be run on a variety of conventional hardware platforms. Data storage 109 contains a corpus of documents, as well as data describing context vectors. Alternatively, the corpus of documents (or other information items) may be remotely located, with electronic links to system 100. In accordance with the software program instructions, CPU 107 accepts input from input device 111, accesses data storage 109, and uses RAM 108 in a conventional manner as a workspace. CPU 107, data storage 109, and program storage 112 operate together to provide a mechanism for generating context vectors and for retrieving information in response to queries.

In the embodiment illustrated herein, CPU 107 can be a mainframe computer or a powerful personal computer; RAM 108 and data storage 109 are conventional RAM, ROM and disk storage devices for the CPU; and output device 110 is a conventional means for either printing retrieved information items, displaying the information on a video screen using a window-based interface system, or sending it to a database for later access.

The preferred embodiment of the present invention is capable of handling a very large corpus, containing over 10 million documents. The architecture supports operation in a distributed data /distributed processor environment, if desired. It may be implemented on any platform, operating system, and user interface of sufficient power and flexibility, such as: XWindows/MOTIF; Sun/OS SunView; Microsoft Windows, VAX/VMS, and the like.

The present invention is based upon a representation scheme using context vectors. A context vector is a fixed length vector having a plurality of component values that are determined

based on relationships between meanings of information items. Such information items may be words, paragraphs, queries, documents, images, and the like. In the following discussion, for illustrative purposes, context vectors are described with reference to words and documents, although many other types of information items may be similarly represented. In the preferred embodiment, each context vector has 200 or more component values.

- The context vector generation scheme of the present invention is designed to produce vectors that represent the relative proximity of meaning or content among words or documents in a quantitative, geometric manner. Thus, information items having similar meanings have closely aligned vectors, while information items having dissimilar meanings have orthogonal vectors.
- This representation scheme allows proximity of meaning to be assessed by performing a simple dot product (inner product) operation on associated context vectors; the higher the dot product result, the more similar the meanings.

Accordingly, the absolute orientation of a particular vector in the vector-space is irrelevant, as long as the relative orientation (with respect to other vectors) is representative of relative proximity of meaning and content. In other words, the problem of finding a set of vectors defining relative meaning has an infinite number of physical solutions in vector-space (absolute orientations) but only one logical solution (relative orientations). The context vector generation scheme of the present invention is designed to arrive at the logical solution through a deterministic training method, without regard to absolute orientation of the vectors themselves.

Context Vector Training

Context vectors are developed for individual words or terms based on proximity to other words. This learning technique is performed on a training set of documents. Referring now to FIG. 1B, there is shown a block diagram of the training system. A training text 101, stop list 103, and phrase list 104 are provided to a preprocessor 102. Training text 101 includes a set of documents for training. Stop list 103 includes a list of words that are deemed uninteresting and are not to be considered in training (e.g., prepositions and common words). Phrase list 104

includes a list of multiple-word phrases that are to be treated as a single word for training purposes (e.g., “world series”, “golden parachute”, “best man”).

Referring now also to FIG. 3, there is shown a flowchart of the training process. The system starts by preprocessing the documents in the training set. Preprocessing consist of several 5 steps, including: 1) removing stop-listed words from the set of training words; 2) consulting phrase list 104 to locate and mark multiple-word phrases that are to be treated as a single word; and 3) reducing words to “stems” in order to increase the effectiveness of the training process—thus, “investments”, “investor”, and “investing” share the stem “invest” and may be treated alike.

The set of word stems generated by preprocessor 102 is fed to learning system 105 which 10 generates a set of stem context vectors 106 according to the method shown in FIG. 3. Each context vector consists of a fixed number of components (200 or more in the preferred embodiment).

Learning system 105 generates stem context vectors as follows. First, initial conditions are assigned 303. In the preferred embodiment, initial conditions are assigned by generating a 15 random context vector for each stem, consisting of components selected by zero-mean, unit-variance Gaussian random number generation. Since the system uses dot products as the measure of relationship strength, mutual orthogonality is a desirable initial condition. This is due to the fact that near-orthogonal vectors will have dot products close to zero. This near-zero dot product corresponds to a weak initial relationship. Assigning a random context vector provides 20 an initial condition that approximates mutual orthogonality. As will be recognized by those skilled in the art, other techniques of assigning initial conditions can be employed.

The system then starts with the first document 304 and proceeds through every document in the training corpus. For each document, it starts at the first word stem 305 and passes through the document, targeting each word stem, one by one. As each stem is targeted, the system 25 applies 306 a learning law to the target. In the preferred embodiment, step 306 involves the following substeps. First, a window is defined, consisting of a fixed number of word stems appearing on either side of the target stem. In the preferred embodiment, the window includes

three stems on each side of the target stem, although the window can be of any size. The stems within the defined window are called neighbors.

Referring now also to FIGs. 2A through 2F, there is shown an example of the window definition for the first few word stems 201 of a sample document. In FIG. 2A, the target stem 5 202 is “Federal” and the neighbor stems 203 are “antitrust”, “law”, and “simple”. Window 204 only includes three neighbor stems 203 because there are no neighbors to the left of target 202 (since target 202 is the first word stem of the document). Next, as shown in FIG. 2B, the target stem 202 is “antitrust”, and the neighbor stems 203 are “Federal”, “law”, “simple”, and “commercial”. FIG.s 2C, 2D, and 2E proceed similarly. FIG. 2F shows the more general case 10 where window 204 includes three neighbors 203 on each side of target 202. (Note that “Sherman Act” is treated as a single word stem; this is defined in phrase list 104, discussed above).

For each target 202, context vectors of neighbors 203 are used to “influence” the context vector of target 202. The relative influence of each neighbor is weighted by two factors: 1) a function dependent on the neighbor’s position in the window relative to the target, and 2) a frequency function determined by the number of documents containing the neighbor stem 15 (frequency). The closer the neighbor, and the lower the frequency function, the more “influence” the neighbor has. These weighting techniques are described below.

The target vector T_j of a target word_j is updated using some sort of learning law. Several different learning laws have been developed, any of which may be used. In some learning laws, a 20 weighted average of the neighbor context vectors is determined and then applied to the target vector T_j ; in other learning laws, the influence of neighbors is determined and applied individually. The most effective learning law has been found to be the following: An error E_{ij} can be defined for each neighbor, representing the difference between the neighbor vector and the target vector. Thus:

$$25 \quad E_{ij} = N_{ij} - T_j \quad (\text{Eq. 1})$$

where:

N_{ij} is the context vector for neighbor of the target word i of target stem j; and

T_j represents the context vector for target j.

A correction value C_j for target j is determined using the formula:

$$C_j = \sum_i^{\text{WS}} (\|E_{ij}\| - \alpha_{ij}) \hat{E}_{ij} \quad (\text{Eq. 2})$$

where:

5 WS is the window size in stems; and

α is a proximity constraint that limits the correcting effect of each error vector E_{ij} on the vector T_j of target_j according to a relationship between neighbor_i and target_j.

If desired, a proximity weight for each neighbor may be applied in generating the correction value C_j. The proximity weight causes neighbor stems that are closer to the target to have a greater effect than those that are farther from the target.

10 The correction value is then applied to the target vector as follows:

$$T_j^{\text{NEW}} = T_j^{\text{OLD}} + \frac{\gamma}{F_j} \sum_{i=1}^{F_j} C_j - M \quad (\text{Eq. 3})$$

where:

15 γ is a predefined step size, or learning rate;

F_j is the total number of occurrences of stem j in the corpus; and

M is the mean context vector for all unique stem vectors in the corpus.

Although the above-described learning law has been found to be highly effective, alternative learning laws are herein presented for illustrative purposes. As will be apparent to those skilled in the art, other learning laws may also be applied.

20 One alternative operates as follows. First a weighted sum W_j of the neighbor vectors for target j is determined using the formula:

$$W_j = \sum_i^{\text{WS}} \frac{G(i)}{D_j} N_{ij} \quad (\text{Eq. 4})$$

where:

$G(i)$ is a Gaussian weight for the neighbor i ; and

D_j is the number of documents that contain target stem j .

Thus, the weighted sum is inversely proportional to the number of documents that contain

5 target stem j . The weighted sum W_j is then added to the target vector as follows:

$$T_j^{NEW} = T_j^{OLD} + W_j \quad (\text{Eq. 5})$$

Another alternative learning law, known as a “batched normalized error” law, operates as follows: First the weighted sum calculated in Eq. 4 is used to determine a weighted error:

$$E_j = W_j - T_j \quad (\text{Eq. 6})$$

A correction value C_j for target j is determined using the formula:

$$C_j = \frac{F_j}{\sum_{i=1}^n E_i} \quad (\text{Eq. 7})$$

Then the correction value is applied to the target vector as follows:

$$T_j^{NEW} = T_j^{OLD} + \gamma C_j - M \quad (\text{Eq. 8})$$

In Eq. 8, the weighted error is not normalized. Alternatively, it may be normalized,

15 yielding the equation:

$$T_j^{NEW} = T_j^{OLD} + \gamma \hat{C}_j - M \quad (\text{Eq. 9})$$

One problem with these learning laws is that they may lead to collapse or convergence of the context vectors (overtraining) when multiple passes are attempted. The solution is to introduce constraints that prevent vectors from getting too close to one another. This was done 20 in Eq. 2. The effect of a constraint is as follows. When the magnitude of the error vector is

greater than α , the target is moved towards its neighbors. When the magnitude of the error vector is less than α , the target is moved away from its neighbors. Thus, convergence and collapse of the context vector set are avoided.

For effective training, α should be selected in accordance with the relative "importance" of the target and its neighbor. If two stems are "important" and they co-occur, then the learned relationship between them should be strong (i.e., the dot product of the two vectors should be large). Therefore α should be small, permitting the two vectors to get closer to each other. In order to vary α according to the importance of both the neighbor and the target, the following technique is employed.

Importance is determined by the frequency of occurrence of the stem in the corpus. The importance metric I_j for stem j varies from a predefined lower bound B to 1.0:

$$B \leq I_j \leq 1.0 \quad (\text{Eq. 10})$$

Importance is determined by the equation:

$$I_j = B + (1 - B) \left(1 - \frac{\log\left(\frac{1}{ND_j}\right)}{\log\left(\frac{1}{TND}\right)} \right) \quad (\text{Eq. 11})$$

where:

ND_j is the number of documents that contain stem j ; and

TND is the total number of documents in the corpus.

From this equation, it can be seen that a stem appearing in every document will have an importance equal to B , while a stem appearing in only one document will have an importance equal to 1.0. The greater the frequency of occurrence, the less important the stem.

The value of α for a particular pair of stems i and j is determined by the equation:

$$\alpha_{ij} = 1 - I_i I_j \quad (\text{Eq. 12})$$

Thus, the greater the co-importance of the two vectors, the smaller the value of α . The value of α is bounded as follows:

$$0 \leq \alpha_{ij} \leq 1 - B^2 \quad (\text{Eq. 13})$$

It can be seen, therefore, that the value of α determines how close any neighbor vector
5 can get to the target vector. The value of α determines the minimum angle between the two vectors, and thus, the maximum dot product between them.

Referring again to FIG. 3, whichever learning law is used, the system then checks 307 to see whether there are more stems to be targeted in the current document. If so, it targets the next stem 308 and returns to step 306.

10 Once the context vectors for all stems of the document have been targeted, the system determines 329 a summary vector for the document as a whole. This vector is representative of the overall content or meaning of the document. It may be generated by simply adding the context vectors of all the stems in the document and normalizing the result. Alternatively, stems may be weighted according to their frequency; the context vectors of stems that occur less
15 frequently in the corpus as a whole are weighted more heavily when the summary vector for a document is calculated. Other techniques of determining a summary vector from a set of stem context vectors may also be employed. Referring now also to FIG. 4, there is shown a block diagram of the process of determining a summary vector for a document according to the preferred embodiment. Context vectors 106 are weighted according to an inverse frequency
20 weight 401 and combined with a vector summing operation 402. The result is normalized 403 to produce a normalized summary vector 404.

Referring again to FIG. 3, if there are more documents to be processed, the system goes to the next document 311 and repeats steps 305 through 309.

Target vectors are not actually updated until the end of an iteration (one pass through the
25 corpus). This prevents flip-flopping and unwanted feedback effects as vectors move towards their neighbors. All corrections are summed during the iteration and applied at iteration

boundaries, in a vector update 312. After all targets T_j are updated, the mean M is also updated 313.

The system then checks 314 whether additional iterations are required for the corpus.

The number of iterations to be performed for a document depends upon some kind of predefined
5 stopping criterion that may incorporate speed, stability, or other concerns. If additional iterations are required, the system returns to step 304.

As will be apparent to those skilled in the art, many variations on the above training techniques are possible, particularly where the training data are non-textual in nature (such as imagery, sound, video, and the like). For illustrative purposes, a more generalized description of
10 a context vector generation technique is given in Development of Context Vectors by Singular Value Decomposition, below.

Context Vector Storage and Indexing

Once the summary vectors have been determined, they are stored. Storage of the normalized summary vectors can be arranged to further reduce searching time by creating cluster
15 trees. An initial parent node at the top of the tree indexed as level 0, node 1, initially contains all of the normalized summary vectors in the data base. A series of child nodes each branching from the initial parent node is created at a next level of the cluster tree. A centroid-consistent clustering algorithm is used to distribute the summary vectors among the series of child nodes. A group of clusters is centroid-consistent if every member of every cluster belongs to the cluster in
20 the group with the closest centroid. A centroid is determined by taking, for each dimension, an average of the component values from all of the context vectors within the group. One popular centroid-consistent clustering algorithm is convergent k-means clustering. Convergent k-means clustering can be performed as follows:

1. Begin with any initial partition that groups the vectors into k clusters. For example, take the first k summary vectors as single element clusters and set the initial value of the centroid of each cluster to equal its member vector. Assign each of the
25

remaining summary vectors to the cluster having the nearest centroid. After each assignment, recompute the centroid for the cluster which gains a vector;

2. Take each summary vector in sequence and compute its distance from the centroid of each of the k-clusters. If the vector is not currently in the cluster with the closest centroid, move the vector to that cluster and update the centroids of the clusters that gain or lose a summary vector;

5 3. Repeat step 2 until convergence is achieved, that is until a pass through all of the summary vectors causes no new assignments.

Since convergence may be rather time-consuming to achieve, the clustering algorithm can

10 be simplified by limiting the number of repetitions of step 2. After a given number of repetitions, such as 99, the centroids can be frozen. Then, one or more passes can be made through all of the summary vectors, distributing the vectors to appropriate clusters, but without updating the centroids. While, using this approximation, the centroids will no longer be exact centroids, the approximate centroids will be sufficient for the use of the present invention. It is not necessary
15 to the present invention that the centroids be precise; however, it is preferable that the clusters be centroid-consistent. The last pass through the summary vectors guarantees that the clusters are centroid-consistent with the approximate centroids. From herein, "centroids" as used in this application shall mean approximate centroids—in other words, a centroid sufficient to establish centroid-consistent clusters. Each node is identified by its centroid for use in the searching
20 process.

Alternatively, other centroid-consistent clustering algorithms may be employed. Such alternatives are described, for example, in "Self-Organized Formation of Topologically Correct Feature Map"; T. Kohonen, Biological Cybernetics, vol. 43, p. 50-69, 1982; and "Performance Evaluation of Self-Organized Map Based on Neural Equalizers in Dynamic Discrete - Signal
25 Detection"; T. Kohonen, et al., in Artificial Neural Networks; and "The Self-Organizing Map", Tuevo Kohonen, Proceeding of the IEEE, Vol. 78, No. 9, September, 1990, all of which are incorporated herein by reference.

In forming a next level of clusters, the nodes in the level above become parent nodes to a set of child nodes below. Only the summary vectors assigned to a parent node are used in the clustering algorithm to form the child nodes which branch from that parent. This is repeated across the entire level of parent nodes and on subsequent levels so that fewer and fewer context
5 vectors are assigned to the child nodes on each lower level. The nodes form a tree pattern in which each node branches from a node in the level above. Thus, each summary vector is assigned to a node on each level of the cluster tree. Each node has a centroid. The bottom-level node assignments for each summary vector and the centroids for each node are stored for use in the search and retrieval algorithms. On the bottom level of the tree, each node points to each
10 normalized summary vector assigned to it. The nodes on the bottom level may be referred to as buckets.

Once a cluster tree has been set up, it is a simple matter to add a new record summary vector to the tree. The initial branches of the tree are examined to find the closest centroid. The summary vector is assigned to the node with the closest centroid. Then, the branches from that
15 node are examined for the closest child node centroid, and the process is continued until a bucket is reached. The new record is then assigned to the bucket with the closest centroid of those buckets branching from the node on the previous level to which the summary vector was assigned. The centroids themselves are not changed. This action preserves centroid-consistency of the clusters. If a bucket gets too large, the summary vectors on the bucket can be divided into
20 subclusters on a subsequent level.

Retrieval

Referring now to FIG. 10, retrieval of context vectors that have been stored according to the above-described tree technique proceeds as follows. The user makes an inquiry using a set of words or terms, or by specifying one or more records (or documents) for which similar records
25 are sought. The words, terms, and records may be weighted if desired to designate which are

most pertinent to the results being sought. After appropriate weighting, context vectors for the words, terms, and records are combined by addition to produce a single query vector 1002.

The query vector is then compared with each summary vector in the data base by performing a dot product vector operation 1003. Lists of these results are accumulated 1004, and

5 the summary vectors resulting in the highest dot products are considered most relevant to the search. If desired, documents can be ranked 1005 in order of dot product magnitude to form a list indicating relevance. The ranked list is then output to the user 1006. Rank ordering by proximity prevents information overload to the user (unlike conventional Boolean search methods, where search results may include, for example, 500 documents, without any indication

10 of which documents are likely to be the most relevant).

By using the above-described cluster tree storage mechanism, the searching task can be greatly accelerated. The query vector is used to identify the summary vector that is closest to the query vector (by dot product computation). The search is performed using a depth first tree walk. A branch is followed down the tree, taking the node at each level having the centroid closest to

15 the query vector. The search proceeds down the tree until a bottom level node (bucket) without children is reached. Each of the summary vectors in the bucket is compared with the query vector (again by dot product computation) to identify the closest summary vector.

Before a subsequent node in the depth first tree walk is checked for a closest vector, first it is determined whether the node can be completely pruned. A node is pruned if it is not

20 possible for a closer normalized summary vector to be assigned to the node than the closest normalized summary vector found so far without violating centroid-consistency. Suppose we are examining a node with centroid C' for pruning. If C is the centroid of any sibling node, then if it is true that any vector closer to the query vector Q than V (closest vector found so far) must be closer to C than C' , then we may prune the node with centroid C' as well as any nodes branching

25 therefrom. This may be computed by comparing the distance between C and C' with twice the sum of the distance between C and Q and the distance between Q and V . If the distance between C and C' is greater, then the node with centroid C' (and its descendants) may be pruned. If not,

the formula is repeated for the remaining sibling nodes, since any one of them may permit pruning to proceed. If none of the sibling nodes achieve pruning of the node, then the search continues through the node with centroid C' and down into the subsequent level if there is one. By using the pruning formula, a node can be pruned when any vector closer to the query vector than the closest vector V must be closer to the centroid C than to the centroid C'. Therefore, that vector could not be assigned to node C or else it would violate centroid-consistency. If this is a bottom node, then all of the summary vectors on the node must be checked to determine whether any are closer than the closest vector found so far. If a closer summary vector is found, it will then become the closest summary vector being remembered. Thus, bottom nodes are thoroughly searched if not pruned. The search continues in a depth first tree walk, pruning off entire branches when possible. These searches continue through the tree until all branches have either been checked or pruned. After the entire tree has been searched, the closest summary vector has been identified. The record associated with the summary vector can be retrieved.

Other known techniques for node pruning, including linear programming techniques, may also be employed.

The system may also employ relevance feedback, whereby the user specifies which of the retrieved documents are most helpful. A new search may then be performed using the summary vector for the specified documents. This technique reduces the time required for searches and improves system effectiveness.

In addition, the most relevant portions of retrieved records may be highlighted if desired. This is done by dividing each retrieved records into a number of sections, representing chapters, paragraphs, or other components. A summary vector is generated for each section, based on the word stems in that section. Dot product computation of the section summary vectors with the query vector is then performed to isolate those sections that are most relevant to the query. The selected sections are then displayed using some distinguishing visual attribute (bold, larger type, a different font or color, an enclosing box, and the like). Thus, the user is able to quickly locate the portions of the document that are most relevant to the query.

One of the possible applications of the above-described system is in the area of conventional ICD9 codes that are commonly used to describe medical procedures. For example, context vectors could be developed to represent medical procedures and their associated ICD9 codes. Then, when additional information is needed for a medical procedure, a query vector

5 could be formulated to retrieve procedures and codes that are relevant to the current procedure.

Another application of the information retrieval system described above is the automated coding of text documents according to a defined index of terms. For example, the Wall Street Journal uses an index of approximately 150 terms to code each article. These terms are assigned by human editors. The information retrieval system described above can be used to emulate the

10 performance of the human editor in assigning such index terms, in the following manner:

1. Build context vectors for words using a sample of text.
2. Collect a set of documents that have been indexed by human "experts" (e.g., editors in the case of the Wall Street Journal), called the *indexed collection*, and generate context vectors for these documents.
3. Generate a context vector for the new document to be automatically indexed.
4. Compare the context vector of the new document with the context vectors of all documents in the indexed collection, and identify the best matches (perhaps the 10 best matches).
5. Produce a list of the index terms of each of the best matches and assign a weight to each term that is proportional to the degree of match such that better matching indexed documents have larger weights than indexed documents that have do not match as well.
6. For each unique index term, generate the index term score by adding the weights of each occurrence of that index term in each of the best matching index documents.

7. Sort the list of unique index terms according to the index term score and assign to the new document those index terms at the top of the list.

See HNC MatchPlus Functional Specification section below, which provides a functional specification, including module and data format descriptions, for a preferred context vector generation, storage, and retrieval system according to the present invention.

Document Visualization

Another useful application of stored summary vectors is in the area of visualization of document content. Context vectors provide a mechanism by which meaning and content of documents can be represented in a visual form, allowing a human observer to take advantage of visually-oriented pattern recognition skills to find documents that are of interest.

Once vectors have been established using the above-described methods, they can be represented visually using any of a number of techniques. The preferred embodiment provides a color graphics visual representation of a set of documents on a computer screen, such as a color graphics workstation or a PC or PS/2 computer equipped with a graphics board. It operates using software written in the C programming language and runs under the UNIX operating system. Essentially, the summary vectors for documents and other information items are displayed in a pseudo-orthogonal display having axes corresponding to query terms (the query terms need not actually be orthogonal to one another). In addition, the items may be displayed with visual attributes representing relative similarity of meaning with other query terms.

Referring now to FIG. 11, there is shown a flowchart of the method of document visualization of the present invention. The method is described with regard to documents, although it may be applied to any type of information items. First, the system accepts 1102 a query from the user. This query may be in the form of a word, term, phrase, document, or other form, as discussed above. Then, the query is separated 1103 into a number of components. This separation may be performed manually by the user, or some automated means may be used, to generate components that maximize information content for the display. For example, query

components may be selected in alignment with the principal components of the document set covariance matrix. These are obtained by considering the data object's context vectors as points in feature space. These points form a cloud, with one point for each object. Principal component analysis fits the best-fitting ellipsoid to this cloud, based on root-mean-squared analysis. Query 5 terms corresponding to the longest perpendicular principal axes of this ellipsoid are selected as the principal components of the data set.

For each component, a display coordinate (axis) or visual attribute is assigned 1104. For example, if five components are identified, the first three may be assigned to the X,Y, and Z axes, and the remaining two may be assigned to the visual attributes of color and texture. Any 10 number of axes and attributes (including visual and non-visual attributes) may be identified and assigned to components, although an excess of axes or attributes may lead to a confusing display.

A context vector is determined for each query component 1105 by the method previously described. Then, for each document to be represented in the display, a set of dot products is computed 1106 using the summary vector of the document with the context vector for each query 15 component. The resulting set of dot product results for each document specifies coordinates and visual attributes for a representation of that document. Once such information has been developed for all documents to be displayed, the results are supplied 1107 to a display engine capable of on-screen display of the icons. Thus, a set of icons are shown on a display screen, having on-screen positions and attributes corresponding to the dot product results.

The display engine may employ known artificial reality image generation technology to portray each document as a three-dimensional icon with a specific shape, size, color, texture and movement projected into a higher dimension context vector space, in accordance with the dot product results previously determined. In addition, coded information about a document (such as the author or the date of publication) can also be represented. 20

Many variations on the above-described visualization scheme are possible. If desired, the above technique may be employed using only display coordinates; all icons will then be displayed having uniform visual attributes. Alternatively, icons having visual attributes such as 25

color, size, and the like, could be displayed without using a positioning scheme. Thus, a relatively compact display could be generated, without a need for high-powered processors to generate the artificial reality display described above. In one embodiment, each icon contains one or more small images of thermometers, each thermometer indicating the degree of

5 correlation (dot product result) with a particular concept.

The display engine of the preferred embodiment is a high-level graphics software interface such as the Programmer's Hierarchical Interactive Graphics System (PHIGS). Other display engines may be used as well. PHIGS and other systems are described below, as well as in the following publications which are incorporated herein by reference: Hill, F.S., Computer
10 Graphics, Macmillan, New York, 1990; Kessener, L.R.A., Data Structures for Raster Graphics, Springer-Verlag, Berlin, 1985; Foley, J.D., and van Dam, Fundamentals of Interactive Computer
Graphics, Addison-Wesley, Reading, MA, 1983.

The description of a three-dimensional scene used as input to PHIGS is simply the definition of each individual object to be displayed, expressed as a set of linked polygons located
15 in a fixed 3-dimensional coordinate system, with each polygon having specified light reflectivity properties (color, specularity, texture, etc.). The polygons make up the visible exterior surfaces of the objects to be displayed. PHIGS handles the lighting of the objects and the calculation of their appearance to the user from a particular vantage point.

The use of such visual display techniques allows a user to view large groups of
20 documents simultaneously in a multi-attribute space. The display of the present invention simultaneously shows the user all of the attributes of each data object for a large set of data objects.

Referring now to FIG. 5, there is shown a set of examples of six different visual attributes
501-506 for icons: size 501, shape 502, color 503, distortion 504, orientation 505, and motion
25 506. Many other examples are possible, including additional visual characteristics as well as sound.

Referring now to FIG. 6, there is shown a sample display having three axes 601, corresponding to the terms "money laundering," "drugs," and "terrorist." Referring now to FIG. 9A, there is shown a sample display with axes 601 and clusters of icons 901 positioned in the coordinate space defined by axes 601.

5 Referring now to FIG. 7, there is shown an example of icon display. In this example, the user has specified two icon words: "Sendero Luminoso", having the attribute of size; and "Simon Bolivar", having the attribute of distortion. Thus, large icons as shown in box 701 indicate a strong association with Sendero Luminoso, while small icons as shown in box 702 indicate a weak association with Sendero Luminoso. Similarly, distorted icons as shown in box
10 703 indicate a strong association with Simon Bolivar, while undistorted icons as shown in box 704 indicate a weak association with Simon Bolivar.

In the present invention, each individually resolvable icon is portrayed as an easily identified object in its proper position, and possessing its assigned attributes. Icons are displayed in simulated three-dimensional space, perspective, and hidden line removal. By means of simple
15 mouse commands, the user is able to navigate through the three-dimensional projection of the higher dimensional context vector space. A user-selected window is available to show the entire vector space, as well as the position and orientation of the user's current viewpoint. The position and orientation can be changed in response to user commands. These operations are performed using conventional computer graphics and artificial reality techniques.

20 Referring now to FIG. 8, there is shown an example of user navigation. Starting view 801 includes three objects 804 positioned with respect to three axes 805. The user selects two operations 803, a rotation and a translation, resulting in ending view 802. In ending view 802, axes 805 have moved, and objects 804 have been repositioned accordingly.

When an icon is located at too great a distance from the user's position to accurately
25 represent all of its characteristics, it will be shown as a point of light. If a large number of icons are located close to one another, they may be shown as a cloud or a shaded region.

If desired, the user may specify data object qualification parameters to help reduce visual clutter and information overload. One method of qualification is to allow the user to specify Boolean parameters, and only display icons which match the specified parameters. Alternatively, the display may be limited to the top-ranked documents resulting from a context vector query 5 performed as described above. Referring now to FIG. 9B, there is shown an example containing axes 601 and one cluster 901 of icons that has been qualified. Other icons outside cluster 901 are not displayed.

Another way to reduce information overload is to provide hierarchical organization of the icons. The user selects an icon for examination of subordinate objects, and specifies a “zoom in” 10 command. When the system zooms in on an icon, all icons representing other documents are erased from the display. New “sub-icons” are introduced, representing sections, chapters, and/or paragraphs of the selected document. These sub-icons are displayed in the same manner as icons.

Responsive to some command, such as a double-click, associated with one of the icons or 15 sub-icons, the associated document (or portion of a document) may be displayed for perusal by the user. Referring now to FIG. 9C, there is shown an example of a text window 910 superimposed on a display having axes 601 and an icon 909. Window 910 contains the text of the document associated with icon 909. The user can move or scroll, adjust window size, and close the window as desired, using conventional window manipulation techniques.

20 ****

Development of Context Vectors™ by Singular Value Decomposition

A NEW APPROACH TO BUILDING MUTUAL SIMILARITY OF USAGE RELATIONSHIP REPRESENTATIONS

1. Introduction

25 Context vectors are a new and powerful approach to the representation of mutual similarity of usage (and/or meaning) between objects in a large data corpus. Context vectors are

being used for document retrieval and routing and for visualization of the contents of groups of documents.

Beyond these proven uses, it is also believed that context vectors can be used for retrieval and routing of foreign language databases and for the construction of content-addressable image, sound, and video databases. Applications of context vectors to the detection and localization of interactions between outlaws (e.g., international criminals, gangs, and terrorists) are also being considered.

The technical approach discussed here builds upon a number of existing concepts, ideas, and results. We have combined and improved upon these important past works to achieve something qualitatively greater. However, at least two critical elements of our approach are entirely new--namely, the idea of context vectors and the observation that proximate co-occurrence is cognitively central to human mental association, and therefore applicable to the assessment of similarity of usage or meaning for almost all raw data--text, imagery, sound, video, etc.--see Section 2 for details.

This White Paper addresses a new possible method for generating context vectors much more efficiently and more accurately than the current method allows. This new method employs principal component analysis, implemented using Singular Valued Decomposition (SVD) and other related techniques.

To provide background for the discussion to follow, the context vector concept is defined and explained in the next section. Following this, Sections 3 through 6 discuss the proposed SVD core context vector generation method and its application to text retrieval and routing. Section 7 discusses the use of context vectors in building a multi-language text retrieval and routing system.

2. Context Vectors

In this section the concept of a set of context vectors is defined. This definition is general and can be applied to text, imagery, sound, and video.

The fundamental idea behind the context vector concept is the observation that many large bodies of real-world data have the property that the statistics of proximal co-occurrence of their basic elements are fixed. Further, frequent proximal co-occurrence of two basic elements can, in almost all cases, be interpreted as implying a strong similarity of usage or "meaning" 5 between these elements. For example, in all human languages, in both spoken and written form (here basic elements are words), elements that frequently appear near one another have a strong associational linkage that could be called "similarity of usage or meaning". Similarly, still or video image elements that frequently appear near one another (with "near" taking on a spatiotemporal meaning in the case of video) can also be taken to have a similarity of usage or 10 meaning. The utility of these observations for data retrieval and routing probably stems from the fact that the associations formed by the human mind are themselves derived, at least in large measure, by training processes driven by common proximal co-occurrence (although, of course, many other things are learned about such data as well). In effect, proximal co-occurrence can be viewed as a first-order approximation to the associational processing of the human mind. What 15 has been discovered during the TIPSTER Program work is that this first-order approximation may be more than sufficient for carrying out high-performance retrieval and routing.

Context vectors are the quantifications of the statistics of proximal co-occurrence. We now define them.

We assume the existence of a large data set (text, imagery, sound, video, etc.). By large it 20 is meant that if the statistics of the structures we will study are determined using a (randomly chosen) sizable subset of the database that these statistics will not, with very high probability, change significantly if re-measured on the whole data set.

We assume that our data set is densely endowed with what we will call elements, each of which belongs to one of N classes. For example, in text, the elements are words that belong to a 25 designated lexicon (other words are ignored). Each lexicon word in the database belongs to one of N classes (stems). In imagery, the elements might be the objects designated by an automated attentional focusing system. These elements would each be assigned to a single class (where the

classes are self-defined as distinct "clusters" of range/azimuth/elevation/background--insensitive feature vectors).

For convenience, we will assume that each element A in the database is numbered with a unique integer index i. We will refer to the i^{th} element of the database as A_i . The number of the 5 class to which A_i belongs will be denoted by c_i , where $1 \leq i \leq N$. The set of indices of all elements belonging to class K, $1 \leq K \leq N$, will be denoted by S_K .

Each time an element appears in the database other elements typically appear "near" it.

For example, when the word "context" appears in this document the word "vector" often appears next to it. It is these proximate co-occurrences that we exploit. To carry out this exploitation we

10 must be given co-occurrence examples. Each such co-occurrence example consists of a base element A_i along with a set of other elements found in close proximity to A_i . For each such proximal element A_j in a co-occurrence example with base element A_i a proximity weighting d_{ij} is defined. The proximity weighting lies between 0 and 1, with greater weights representing greater proximity or nearness. Note that each database element can appear as the base element in 15 at most one co-occurrence example. Some elements may not appear as base element any co-occurrence examples because we might not have all possible co-occurrence examples available to us. Finally, if we are given a co-occurrence example with base element A_i in which A_j appears with proximity weighting d_{ij} , then we assume that we will also be given a co-occurrence example with base element A_j in which A_i appears with the same proximity weighting $d_{ji} = d_{ij}$. In other 20 words, we assume that proximity weightings are symmetric.

Given a large set of co-occurrence examples, we then form the square mutual co-occurrence matrix $R=(r_{uv})$, where:

$$r_{uv} = \text{AVERAGE}[d_{ij} | d_{ij} \neq 0 \text{ and } i \in S_u \text{ and } j \in S_v],$$

with the averages being taken over all co-occurrence examples available for training the system.

25 Note that since we have assumed that for each d_{ij} we have an equal d_{ji} in the training examples, R is automatically symmetric. Also note that since d_{ii} can be assumed to be 1 (each element co-

occurs with itself with proximity weighting 1), that R has its diagonal elements all equal to 1. To support the work that follows we note that because R will typically be very sparse and because almost all of the off-diagonal elements will be less than 1, it is not unreasonable to make the following assumption:

5 Assumption 1: The matrix R is assumed to be positive definite.

As shown by Strang [20], this is equivalent to saying that the determinants of the upper-left, diagonal-centered square submatrices of R all have positive determinants. This can normally be expected to be true in the cases of text, sound, imagery, and video (and can be tested for any specific R).

10 Given the above preliminaries, we are now ready to define context vectors.

Definition 1: A set of vectors

$$\{w_1, w_2, \dots, w_N\} \subset \mathbb{R}^n,$$

where n is an integer with $n \gg 1$, is called a set of context vectors for a data set mutual co-occurrence matrix R iff

15 $w_i \cdot w_j = r_{ij}$

for all $i, j = 1, 2, \dots, N$.

The meaning of this definition is that the context vectors (which are all unit-length vectors, since $w_i \cdot w_i = r_{ii} = 1$) have directions that represent the mutual co-occurrences of the element classes. Note that since the w_i are unit vectors that $w_i \cdot w_j = r_{ij} = \cos(\theta_{ij})$, where θ_{ij} is the angle between w_i and w_j . As mentioned above, the context vectors encode the relative usage or, in some crude sense, the relative meaning of the element classes. Note that if we were to rigidly rotate the entire context vector set in \mathbb{R}^n that we would get an equivalent set of context vectors (since it is only the mutual angles between the vectors that matter). Thus, context vectors are not unique.

25 Now that we have precisely defined what a context vector is, some questions arise. First and foremost, given a mutual co-occurrence matrix R, do a set of context vectors exist? Another

question is: given R , can we create a set of context vectors from it? Finally, how small can n be?

These are some of the questions we examine in the next two sections.

3. Context Vectors Exist And Can Be Constructed

Given a mutual co-occurrence matrix R it is easy to show that a set of context vectors

5 must exist for it.

Theorem 1: Given a symmetric, positive definite mutual co-occurrence matrix R , there exists a set of context vectors in R^N for it.

Proof: If we define the $N \times N$ matrix W to be

$$W = [w_1, w_2, \dots, w_N],$$

10 where the $N \times 1$ w_K vectors are the columns of W , then the condition that the context vectors must meet (from Definition 1) can be re-expressed as

$$W^T W = R.$$

Thus, we must show that there exists such a matrix W . To do this all we need note is that since R is symmetric and positive definite we can use Gauss decomposition to reexpress this 15 matrix as

$$R = QDQ^T,$$

where the columns of the orthogonal matrix Q are the unit-length eigenvectors of R and the matrix D is diagonal, with the eigenvalues of R as its diagonal entries (see Strang [20] for details). If we then take the square root matrix of D (namely, the matrix $D^{1/2}$ which has each of 20 its diagonal elements equal to the square root of the corresponding diagonal element of D) and rearrange terms we get

$$R = QDQ^T = QD^{1/2}D^{1/2}Q^T = (D^{1/2}Q^T)^T (D^{1/2}Q^T).$$

Thus, we can take W to be $D^{1/2}Q^T$. Thus, for any symmetric, positive definite R there exists a W in R^N .

25 The upshot of Theorem 1 is that if we are willing to let $n=N$ (which, for most applications, will be a very high dimensional space), then we can always find a set of context vectors in R^N , no matter what R is. Namely, all we need to do is calculate the eigenvalues and

eigenvectors of R and then use them to construct $W=D^{1/2}Q^T$. While this indeed gives us a set of context vectors, they are not really ones that we want. Constructing desirable sets of context vectors in R^N , where $n < N$ is the subject of the next section.

4. SVD Construction of Practical Context Vectors

5 While the formula $W=D^{1/2}Q^T$ for constructing context vectors will work, it essentially leaves us stuck with using N dimensions for representing these vectors. For a practical application this would typically be a debilitating disadvantage in terms of computer memory and computational burden. For example, in a text data access system we might have $N=250,000$ (the number of stems in the corpus). In this section we demonstrate a method for producing context 10 vectors in a (typically) much lower dimensional space that can perform approximately as well as "full-dimensional" context vectors. In fact, we provide a formula for calculating the error induced by reducing the context vectors to n dimensions from N dimensions.

The basic idea is to employ the Singular Value Decomposition (SVD--see Strang [20] for details) in which any arbitrary MXN matrix S can be written as

15 $S=PD^{1/2}Q^T$,

where P is an MXM orthogonal matrix (with columns given by the unit-length eigenvectors of S^TS), Q is an NXN orthogonal matrix (with columns given by the unit-length eigenvectors of SS^T), and $D^{1/2}$ is an MXN matrix with the square roots of the r non-zero eigenvalues of S^TS (and SS^T --they have the same non-zero eigenvalues, all of which are positive and real) as the entries 20 in its upper left "diagonal" and with zeros in all other positions. The orders of the eigenvectors in P and Q must match the order of the corresponding eigenvalues in D. Note that the SVD is not unique (for example, we can rearrange the orders of the eigenvalues and eigenvectors in D, P, and Q). From now on we will arrange D so that the eigenvalues are listed from largest to smallest, starting with the largest eigenvalue at the upper left corner of the matrix and moving 25 down the diagonal to the right

In our case we will be interested in applying the SVD to yield a lower-dimensional set of context vectors from our matrix R. To do this, we first note that the SVD expansion depends only

on the properties of the matrices $S^T S$ and SS^T . Since we want $W^T W = R$, and since R is symmetric, we can identify R with both $S^T S$ and SS^T . Thus, the eigenvectors of $S^T S$ and SS^T will be the same. So, in this SVD case, we will have $P=Q$. Thus, we can write W as

$$W = Q D^{1/2} Q^T,$$

- 5 with Q and D taking on the same meaning as in the previous section. It may seem odd that this construction of a W is so close to, and yet clearly different from, that provided by Theorem 1. However, as we noted above, W is by no means unique. To see that this SVD-derived W is indeed acceptable, note that

$$W^T W = (Q D^{1/2} Q^T)^T Q D^{1/2} Q^T = Q D^{1/2} Q^T Q D^{1/2} Q^T = Q D Q^T = R,$$

- 10 where we have exploited the property that $Q^T Q = Q Q^T = I$ (the identity) for any orthogonal matrix.

The value of using $W = Q D^{1/2} Q^T$ instead of $W = D^{1/2} Q^T$ becomes clear when we expand this new W in terms of the columns of Q . To make the notation clearer, we let

$$Q = [u_{ij}] = [u_1, u_2, \dots, u_N],$$

where the $N \times 1$ vectors u_k (the unit-length eigenvectors of R) are the columns of Q . Using

- 15 $W = Q D^{1/2} Q^T$ we can then write

$$w_i = \sum_{j=1}^N (\lambda_j^{1/2} u_{ij}) u_j = \sum_{j=1}^N \alpha_{ij} u_j,$$

where the (necessarily non-negative and real) eigenvalues $\lambda_1, \lambda_2, \dots$ are listed in descending order.

By virtue of this formula, we can see that each w_i vector is expressed as a weighted sum of a fixed set of N vectors (the u_j). Because of the nature of the matrix R , it is reasonable to expect

- 20 that many of its eigenvalues λ_j will be close to zero. This is true of almost all real-world matrices constructed in a like manner, and we certainly expect it to be true in this case (see below for some experimental results supporting this supposition). Because of this, and because the u_{ij} values are all small (recall that Q is orthogonal, and thus the u_{ij} are the components of unit vectors), the scalar coefficients multiplying the u_j vectors in the w_i sum become progressively smaller as j increases (regardless of which particular w_i vector we are considering). Thus, it may

be reasonable to truncate these sums after r terms, where $1 < r < N$. In this case, we get the approximation

$$w_i \approx \sum_{j=1}^r \alpha_{ij} u_j.$$

Note that the u_j vectors are an orthonormal basis for \mathbb{R}^N . However, we have now

- 5 discarded all but r of them. So this is an orthogonal basis for \mathbb{R}^r . Thus, we can think of our w_i vectors as belonging to an r -dimensional Euclidean space. If we reexpress our w_i vectors in terms of the u_j basis vectors (i.e., taking the first coordinate of this new representation to be the u_1 component, the second to be the u_2 component, and so on) we get

$$w_i = (\alpha_{i1}, \alpha_{i2}, \dots, \alpha_{ir})^T \equiv v_i.$$

The error introduced by using these r -dimensional approximations v_i can be quantified. Specifically, the errors in the inner products will be given by

$$|w_i - w_j - v_i \cdot v_j| = \sum_{k=r+1}^N \alpha_{ik} \alpha_{jk}.$$

- In conclusion, we have shown that a set of lower-dimensional context vectors can be
 15 constructed by applying the SVD to the mutual co-occurrence matrix R . The error introduced into the context vectors by reducing the dimension from $n=N$ to $n=r$ is quantifiable and can be controlled as required by the specific problem.

5. Applicability of the SVD Method

- The first question that naturally arises in connection with the SVD method proposed
 20 above is how rapidly the eigenvalues of the co-occurrence matrix R actually fall off in a real-world situation. This will determine both the practical dimensionality of the context vectors and the utility of the SVD method. This question is at least partially answered by FIGS. 12 and 13. These figures were created by calculating the eigenvalues of the matrix

$$T=WW^T,$$

where, in this case, the W matrix is defined by a set of context vectors obtained using our current random-initialization and bootstrapping context vector generation method (the columns of W are the context vectors). As is well known, the matrices $W W^T$ and $W^T W$ have the same non-zero eigenvalues. However, $W^T W$ is an $N \times N$ matrix (where N is the number of stems--approximately 15,000 in these examples), whereas $W W^T$ is an $n \times n$ matrix, where n is the dimension of the context vectors (n equals 280 and 512 in the examples of FIGS. 12 and 13, respectively).

Therefore, we calculated $T=WW^T$.

FIG. 12 Ranked eigenvalues (largest on the left to smallest on the right) for $T=WW^T$

derived from a set of approximately 15,000 280-dimensional context vectors (i.e., $N \approx 15,000$ and $n=280$). Note that the smallest eigenvalues are not much smaller than the largest eigenvalues. This suggests that the dimensionality of this context vector space cannot be reduced further (perhaps increasing the dimensionality would improve somewhat).

FIG. 13 Ranked eigenvalues (largest on the left to smallest on the right) for $T=W W^T$

derived from a set of approximately 15,000 512-dimensional context vectors (i.e., $N \approx 15,000$ and $n=512$). Note that, unlike the case of FIG. 12, the smallest eigenvalues are much smaller than the largest eigenvalues. This suggests that the context vectors for this case could be adequately represented in a lower-dimensional space.

To see the effect of changing n on the eigenvalues of T we generated two sets of context vectors: one set for $n=280$ and one set with $n=512$. Once the context vectors were built (using a 1,000 document test corpus of Wall Street Journal articles from the TIPSTER collection) we then formed the matrix T, calculated its eigenvalues, and plotted them from largest to smallest. The results for $n=280$ and $n=512$ are shown in FIGS. 12 and 13, respectively. Clearly, for $n=512$ we see that the space can be reduced somewhat in dimensionality using the SVD method. Thus, we anticipate that the SVD method will be of use in developing sets of context vectors with the smallest possible dimensionality, without reducing the dimensionality too much (as in the $n=280$

case of FIG. 12. For most text data access applications we believe that context vectors with dimensionalities between 400 and 1000 will be adequate.

We now consider the value of the SVD method for applications of context vectors.

6. Practical Context Vector Generation Using SVD

The main incentive for developing an SVD method for context vector generation is that it will enable us to rapidly build highly accurate context vectors for a core set of element classes. These context vectors can then be used with a two-pass bootstrapping method to build the context vectors for all remaining element classes. By core set it is meant that there often exists a set of high-frequency-of-appearance elements that can be used as the foundation for the context vectors for the remainder of the classes. For example, in text we might choose, say 2,000 high-frequency stems to use as a core class set. We would then use a very large corpus to compute the mutual co-occurrence matrix for this core class set (R would be a $2,000 \times 2,000$ matrix). The SVD method would then be used to create a reduced-dimensionality set of context vectors from this R matrix.

The benefit of this method over the current initial random context vector/bootstrapping method is that only one pass through the data set would be required to build an excellent set of core context vectors. Further, since the co-occurrence matrix need only concentrate on the specific classes found in the core set, the number of calculations required during the one pass through the data is greatly reduced (as opposed to the current method, where all of the class context vectors must be adjusted during each of multiple passes through the data). This will allow the SVD method to use a much larger data set in the construction of the R matrix--thus yielding more accurate context vectors.

Once the R matrix is formed, it will be necessary to apply the SVD method to it. Since R will be a large, sparse, symmetric, positive semi-definite matrix we will apply special numerical techniques for this process. A number of these have been developed. For example, Professor Martin Schultz of Yale University has developed a large library of software subroutines for this purpose (they are sold by Scientific Computing Associates, Inc.). His software has been used to

calculate the eigenvalues and eigenvectors for matrices of the above kind with dimensions as high as 100,000X100,000.

Once the eigenvalues of R have been calculated we will then use a set of rules to determine where to set the dimensionalities of the context vectors. The context vectors for the 5 core classes will then be created.

Another possible approach to the SVD method is to use an adaptive neural network method for developing the eigenvalues and eigenvectors of the R matrix. This method has shown great promise, as it allows calculation of these quantities with typically greatly reduced computational burden, when compared with the algorithmic methods discussed above (see [1, 2, 10 3, 7, 10,11] for details on these neural network SVD methods). The primary reason for this reduction is that the R matrix does not have to be computed. These methods work by passing through the data itself. However, these methods will have to be adapted to this application, as our requirements are quite different from those of other applications that have used these methods (e.g., data compression and grammar analysis).

At least two different approaches to building a complete set of context vectors can be tried using the above core context vector generation procedure. One approach is to simply let the core set include almost all the classes. For example, in a typical corpus of approximately 1 million newspaper articles there would be approximately 40,000 stems which would appear more than 10 times in the corpus. If we restricted this to stems which appear 100 or more times this 15 number would be reduced to roughly 10,000. It may be feasible to directly compute the R matrix for these classes and carry out SVD on it. The remaining stems (about 1 million, approximately 75% of which appear only once each) can then have their context vectors computed by adding up the weighted core context vectors in their co-occurrence sets during a single pass through the corpus. In this case, we believe that the SVD method will provide an excellent set of context 20 vectors that have their mutual geometric relationships arranged exactly as dictated by the data 25 itself.

Another approach to the development of a complete set of context vectors is to use a much smaller core set (e.g., the 2,000 highest frequency stems in the case of the TIPSTER corpus). Once the core context vectors have been created, they will be frozen. These core context vectors will then be used to determine the other class context vectors. Two passes of the batch 5 bootstrapping method will be used. The first bootstrapping pass through the training data will be used to collect weighted sums of core context vectors for each co-occurrence set which contains a sufficient number of core context vectors. Those non-core classes that have sufficient numbers of co-occurrence sets will have their context vectors fixed at the end of this pass. A subsequent second pass will most likely be sufficient to calculate the context vectors of the remainder of the 10 classes. It may be desirable to then go back and complete one more pass to readjust the core and first-pass context vectors based upon the second pass results. This method requires three passes through the data, but lowers the size requirement for the core class set.

Both of these SVD context vector generation methods are expected to be much more efficient and accurate than the current random initialization plus bootstrapping method. 15 Significant improvements in retrieval and routing performance are expected from this new method of context vector generation.

7. Context Vectors for Foreign Language Text

One of the pressing needs in data access for text is to be able to retrieve and route documents in all languages that discuss a particular topic described in a single language query. 20 HNC has devised a specific method for solving this problem using context vectors. This method is described in the this section using the example of a corpus containing documents in both English and Spanish.

The first step is to build (or adopt from another system) a set of context vectors for the English portion of the corpus. A person who is skilled in both languages in then engaged. Their 25 job is to create a list of tie words, which are words that have the exact same meaning in both languages. In informal discussions with persons who know both English and Spanish, English and Russian, English and Chinese, and English and Japanese it is clear that it is easy to produce

tie word lists having hundreds of entries for each of these language pairs. With effort, lists of at least two thousand tie words could probably be produced. We presume that this will be possible for any two human languages. To make this process simple, all tie words would be between English and the foreign language (in the case of this section, Spanish). This would seem

5 reasonable because English is arguably the richest human language and it is also second in worldwide popularity (behind Chinese).

Once the tie words have been selected, their context vectors in English are looked up. These are then transferred to the context vector set for the corresponding words in the foreign language. These tie word context vectors are then frozen. The remainder of the context vectors

10 for the language (a stop list and a stemmer can be used, if desired, but we do not believe that these are all that beneficial) are then created, using, for example, a neural network SVD method that adaptively builds new context vectors using the frozen tie word context vectors as a substrate. The final result is a set of context vectors for the foreign language that lie in the same context space as the English context vectors.

15 Key Point: Context Space is a universal meaning representation domain that we expect will be usable for representing the meaning of essentially all human data items (text in all languages, speech in all languages, imagery, and video).

Once the context vectors for the new foreign language have been built then the documents in that language have context vectors built for them and these are logged into the

20 context space database.

In essence, context space becomes a crude universal method of describing the usage or meaning of a data item. The same space is used for all languages (and, eventually, for video, imagery, and sound as well). By means of queries expressed as vectors in this universal context space we can retrieve data in all languages and media on the basis of its content. When finally

25 built, this will be the ultimate content addressable memory system.

To retrieve or route documents in multiple languages requires no new mechanisms. A query in one of the available languages is first formulated (the machine must be told which

language it is). This query is then converted to a context vector. The context vector is then used to search through the document context vector database to find the closest matches, which are then presented in rank order (Boolean queries can also be used, but the key words will only be used with documents in the language from which they come--which must be specified). The net 5 result is a list of the most relevant documents, independent of language. The user can then select which documents they care to look at and they will be displayed in the text window of the system.

In the case of a English and Spanish system, several advantages combine to make such a system much easier to build than, say, an English and Japanese system. First, significant volumes 10 of Spanish text are available on CD-ROM. Second, the characters used in Spanish already exist within our Match Plus. system. Third, many expert speakers of both English and Spanish are readily available to us. For these reasons, we believe that it will be possible to build an English and Spanish system.

Another potential advantage of having a common context space for all languages is that it 15 will probably also be possible to build a crude gisting system. The idea of this would be that a foreign language document would be displayed in a text windows. The user would call up a gisting feature (by using a keyboard control sequence or by means of a mouse-activated menu selection). The gisting feature would place a window approximately one paragraph long in the text. Next to this window (to the side of the highlighted text in the window) would be a second 20 window containing a selection of English words that have context vectors closely aligned with the aggregate context vector of the material in the foreign language window (which would be computed automatically by the gisting system). The English words in the gisting window would be presented in order of closeness to the context vector of the foreign language window.

Although they would not be formed into sentences, we believe that these words would 25 nonetheless give a very clear idea of the content of the selected passage. The user could then scroll the gisting window up and down the foreign language text to survey its content. Further, the existing Match Plus highlighting system could be used to locate those passages of the text

that are most highly related to the subject mane of the query currently presented in the query window. In the end, we believe that this gisting window feature will, in many cases, obviate the need for translation of what are later recognized as irrelevant documents into English. This is yet another analyst productivity enhancement that we expect to flow from context vector technology.

5 REFERENCES

1. Almeida, L. B., and Silva, F. M., "Adaptive decorrelation", in: Aleksander, I., and Taylor, J. [Eds.] Artificial Neural Networks, 2, 149-156, Elsevier, Amsterdam, 1992.
2. Baldi, P., and Hornik, K., "Neural networks and principal component analysis: Learning from examples without local minima", Neural Networks, 2, 53-58, 1989.
- 10 3. Cottrell, G. W., Munro, P., and Zipser, D., "Learning internal representations from grey-scale images: An example of extensional programming", Proc. Ninth Annual Conf. Cog. Sci. Soc., Seattle, Erlbaum, Hillsdale N.J., 1987.
4. Dumais, S., "LSI meets TREC: A status report", Technical Report, Bell Communications Research, Morristown N.J., 1993.
- 15 5. Gallant, S. I., "A practical approach for representing context and for performing word sense disambiguation using neural networks", Neural Computation, 3, 293-309, 1991.
6. Gallant, S. I., "Context vector representations for document retrieval", Proc. AAAI Workshop on Natural Language Text Retrieval, Anaheim Calif., 15 Jul. 1991.
7. Hecht-Nielsen, R., Neurocomputing, Addison-Wesley, Reading Mass., 1991.
- 20 8. Koll, M. B., "WEIRD: An approach to concept-based information retrieval", Preprint, School of Information Studies, Syracuse Univ., Syracuse N.Y., c. 1979.
9. Landauer, T. K., and Littman, M. L., "Fully automatic cross-language document retrieval using latent semantic indexing", Technical Report, Bell Communications Research, Morristown N.J., c. 1990.
- 25 10. Oja, E., "Principal components, minor components, and linear neural networks", Neural Networks, 5, 927-935, 1992.

11. Oja, E., Ogawa, H., and Wangviwattana, J., "PCA in fully parallel neural networks",
in: Aleksander, I., and Taylor, J. [Eds.] Artificial Neural Networks, 2, 199-202, Elsevier,
Amsterdam, 1992.
12. Ossorio, P. G., "Classification space: a multivariate procedure for automatic
5 document indexing and retrieval", Multivariate Behavioral Research, 479-524, October 1966.
13. Osterberg, M. and Lenz, R., "Unsupervised feature extraction from first principles",
Preprint, Image Processing Laboratory, Linkoping University, Sweden, 1993.
14. Raghavan, V. V., and Wong, S. K. M., "A critical analysis of vector space model for
information retrieval", J. of Amer. Soc. Inf. Retrieval, 37, 279-287, 1986.
- 10 15. Robertson, S. E., "Progress in documentation", J. of Documentation, 33, 126-148,
June 1977.
16. Salton, G., "Developments in automatic text retrieval", Science, 253, 974-980, 30
Aug. 1991.
17. Salton, G., and Buckley, C., "Global text matching for information retrieval", Science,
15 253, 1012-1015, 30 Aug. 1991.
18. Salton, G., Wong, A., and Yang, C. S., "A vector space model for automatic
indexing", Comm. of the ACM, 18, 613-528, November 1975.
19. Sparck-Jones, K., "Search term relevance weighting given little relevance
information", J. of Documentation, 35, 30-48, March 1979.
- 20 25 20. Strang, Gilbert, Linear Algebra And Its Applications, Second Edition, Academic
Press, New York N.Y., 1980.
21. Sutcliffe, R. F. E., "Distributed representations in a text based information retrieval
system: a new way of using the vector space model", Proc. SIGIR'91 Conf., ACM, Chicago Ill.
13-16 Oct. 1991.
22. Van Rijsbergen, C. J., "A theoretical basis for the use of co-occurrence data in
information retrieval", J. of Documentation, 33, 106-119, June 1977.

23. Wong, S. K. M., et al, "On modeling of information retrieval concepts in vector spaces", ACM Trans. on Database Sys., 12, 299-321, June 1987.

DOCUVERSE

5 An Intuitive Visual Representation System for Exploring Large Document Sets

1. Executive Summary

A critical problem faced by analysts is the ever growing volume of written material that is available. It is said that well in excess of 90% of all desired intelligence information is sitting in available documents, waiting to be found and digested. Boolean query based document retrieval and routing systems were, historically, the first attempt to find a way to access large document corpora on the basis of topics of interest. Next-generation meaning-similarity based document retrieval and routing systems (in particular, HNC's TIPSTER program system) are now being developed. These are expected to significantly increase the productivity of analysts in terms of increasing their ability to much more quickly and comprehensively access documents that pertain to a particular area of interest.

Although high-performance document retrieval and routing is of critical importance, it only addresses one aspect of an analyst's job. Namely, to probe document databases for information of a known type. Another, equally important, analyst function is to search for unexpected and unusual developments. To look for new trends and emerging activity patterns.

20 Document routing and retrieval systems based upon searching for specified types of information cannot, by their very nature, be of much use in carrying out this exploratory function. What is needed is a way for analysts to somehow "get into" the universe of documents and roam around--to get to know the structure of the document database and look for any oddities or surprises that present themselves.

25 This White Paper proposes a new type of document database exploration tool--one that is designed to allow analysts to exploit their powerful natural visual pattern recognition skills to

explore the information content of huge numbers of documents simultaneously. It will be the analyst's window into the document universe. We therefore call this concept the DOCUVERSE system.

The goal of the DOCUVERSE system is to support an intuitive, easy to control,

- 5 exploration process whereby aspects of the contents of large numbers of documents can be rapidly assessed. The substrate for this process is a color graphics visual representation of a set of documents on a computer screen (e.g., a color graphics workstation or a PC or PS/2 computer equipped with a graphics board). This representation, which will exploit some of the artificial reality image generation technology developed for use in Hollywood films and flight simulators,
- 10 will portray each document, in the set of documents being considered, as a 3-dimensional object or icon with a specific shape, size, color, texture, and movement located in a three-dimensional cyberspace. Each of these attributes associated with each document corresponds to the document's content's similarity of meaning with one of a set of user-chosen features (such as a body of text indicating a topic of interest). Coded information about a document (such as when it
- 15 was written or the identity of its author) can also be used as a feature. This will allow an analyst to view large groups of documents in a multi-attribute space.

The project discussed in this White Paper will develop a prototype DOCUVERSE system. The system will be written in the C software language and it will run under Unix on an off-the-shelf color graphics workstation. It will not depend upon any other special hardware, but

- 20 will utilize the workstation's color graphics display capability and the workstation's mouse.

This project will also include the procurement and installation of a TIPSTER Software Evaluation System Sun Microsystems workstation system (integrated with the DOCUVERSE system workstation) capable of running the software being developed by the TIPSTER document detection contractors.

25 2. Definition of the Problem

 2.1 The End-User Problem

Effective access to large databases of textual information is a continuing operational problem. Ongoing developments in computer networks, query-based retrieval and routing systems, and electronic mail systems are providing ever increasing access to documents of interest on specific topics. However, analysis of the contents of these documents and exploratory discovery of trends and changes in this content must still be carried out manually. Tools for supporting this crucial area of work simply do not exist. Since these activities represent perhaps half of an analyst's job, the development of computer-based tools in this area is of the highest importance.

2.2 The Technical Problem

Advanced searching and retrieval methods provide the capability to locate a large fraction of the information on a specified topic that exists within a document database. The final product of a system based upon these methods is a prioritized list of documents that are relevant to a specified topic. The analyst must then examine these documents and use their content to reach conclusions. Thus, a "one-document-at-a-time" analysis bottleneck is created, which often limits the analyst's ability to quickly identify trends, changes, etc. What is needed is another way to deal with the contents of a preselected set of documents (e.g., the output of a query-based retrieval and routing system) on a mass basis. Specifically, it is desirable to have a capability for viewing the documents as individual objects floating in a visual cyberspace, with the position and display of each object determined by its ranking or correlation with user-defined document attributes.

The goal is to produce a document set content exploration system that can operate on large document sets and which can exploit natural human visual scene analysis capabilities.

3. The DOCUVERSE System Concept

Human vision can perceive and interpret many dimensions of information if the information is encoded and represented correctly. Context vector representation and high resolution displays are enabling technologies for visualization of textual information. Display of information can be accomplished such that humans can "navigate" through abstract representations of textual databases. That is the purpose of the DOCUVERSE system.

This section begins with a review of the HNC technique of context vectors. The DOCUVERSE System concept is then explained via discussions of its constitutive elements.

3.1 Context Vectors

HNC's TIPSTER project document retrieval and routing system and the DOCUVERSE system proposed here are both based upon the use of context vectors. A context vector is a point on the unit radius sphere in n-dimensional Euclidean space that represents the meaning content of a document (or a segment of a document). The components of the vector are the correlations of the overall document meaning with the meanings of some fixed features. Features include carefully chosen highly descriptive terms that represent the concepts of a small context-free language. Other future values are adaptively learned from the corpus. By means of a mathematical technique (see the HNC TIPSTER Project Proposal and Project Documents for details) a context vector can be assigned to any word, phrase, or document segment. As we have shown on the TIPSTER project, these context vectors provide an accurate guide to similarity of meaning--regardless of the nature of that similarity. This new approach can be contrasted with more conventional similarity of meaning techniques, such as the WordNet system.

Prof. George Miller and his colleagues at Princeton University have developed a type of associational dictionary called WordNet. WordNet consists of several data structures--one for verbs, one for nouns, one for adjectives, etc. The verb structure is a mathematical ordered tree structure generated by the relation "in the manner of". A verb in the tree is connected above another verb if the latter verb action is in the manner of the former. For example, the verb walk lies below, and is connected to, the verb move in the verb tree, because walking is an action in the manner of moving. The noun tree in WordNet is similar to the verb tree, except that the relation is "is a kind of". Thus, animal lies above and is connected to dog, because dogs are a kind of animal.

Experiments carried out under the TIPSTER Program have shown that HNC's context vectors capture the very limited kind of similarity found in WordNet, and much more. Although exhaustive testing has not been carried out, spot checking has shown that pairs of words

WordNet considers close in meaning are also considered close in meaning by the HNC context vector system. However, context vectors go beyond this. For example, because of the very limited relations that are coded in WordNet, no relationship at all would be found for the words drive and automobile (since one is a verb and the other a noun). However, the context vectors for 5 these words are quite close, as are those for block and city, and block and tackle. On the other hand, the context vectors for the words automobile and dog are not close. In fact, since the similarity of context vectors for words flows from the statistics of their close proximity in huge volumes of text, essentially all types of similarity of meaning become automatically incorporated into this meaning representation. This carries over into the context vectors for phrases and 10 document segments, since these are constructed by taking weighted vector sums of word context vectors.

Besides providing an excellent vehicle for the encoding of meaning, the mathematical form of the context vectors can be exploited to develop a fast searching method so that the similarity of meaning of each document in a document set can be determined rapidly. It is also 15 possible to do trimmed searches to find only those documents that have a high correlation of meaning with a particular context vector. These properties flow from the facts that comparing similarity of meaning is carried out by the simple mathematical operation of taking the inner or dot product between the selected context vector and those of the documents to be rated as to similarity. Searches through large document sets for close matches can be carried out by means 20 of hierarchical cluster searches, which are very efficient and fast.

In summary, HNC's context vector approach provides a practical means for representing the meaning of a word, phrase, or document and for comparing the similarity of meaning of multiple documents.

3.2 Data Objects

25 The documents that make up the databases of interest for analysis contain information that is hierarchically organized into sets of nested data objects (see FIG. 14) For the purposes of

this proposal, these data objects could be any of the following (other possibilities, such as images, line drawings, and logos will be ignored here):

Words.

Part of a document (sentence, paragraph, section or chapter).

5 Document.

Set of documents.

Database of documents.

Set of databases.

The DOCUVERSE system will be able to display content information on most or all of

10 these types of objects. However, to keep the presentation simple, we will concentrate here only on the case where the data objects being portrayed are all documents.

3.3 Artificial Reality Display Systems

The DOCUVERSE system will employ state-of-the-art computer graphics techniques to create a visual display of the document universe to be explored. This display will be constructed 15 using artificial reality software that will be adapted for this project from an existing software package. The basic ideas behind this software are described in this subsection.

During the 1980s, two very similar standard high-level graphics software interfaces were developed--GKS (the Graphical Kernel System) and PHIGS (the Programmer's Hierarchical Interactive Graphics System). These are now merged into a standard PHIGS interface that is 20 supported by multiple workstation vendors (e.g., Sun Microsystems, Hewlett-Packard, IBM, and Silicon Graphics), along with many enhancements (which we shall automatically include when we hereinafter say "PHIGS"). For details on computer graphics, GKS, PHIGS, X-Windows, and related issues see: Hill, F. S., Computer Graphics, Macmillan, New York, 1990; Kessener, L. R. A., Data Structures for Raster Graphics, Springer-Verlag, Berlin, 1985; and Foley, J. D., and van 25 Dam, Fundamentals of Interactive Computer Graphics, Addison-Wesley, Reading, Mass., 1983.

The basic idea of PHIGS is that the graphics creation job is broken into two pieces: describing the objects to be displayed to the user, and generating the display of those objects

from a user-selected eyeball position, direction of look, and panoramic angle of view (i.e., zoom level). The first job is that taken on by the application software. The second job is automatically carried out by PHIGS.

The description of a 3-dimensional scene used as input to PHIGS is simply the definition
5 of each individual object to be displayed, expressed as a set of linked polygons located in a fixed 3-dimensional coordinate system, with each polygon having specified light reflectivity properties (color, specularity, texture, etc.). The polygons make up the visible exterior surfaces of the objects to be displayed. The lighting of the objects and the calculation of their appearance to the user from his or her selected eyeball vantage point is the job of PHIGS.

10 On some workstations, PHIGS is implemented primarily in software. On others, PHIGS is implemented primarily in hardware. While the ultimate screen appearance is essentially the same in both cases, those workstations that implement PHIGS in hardware are usually able to display imagery at speeds that would be consistent with a human analyst's work pace, whereas those that implement PHIGS primarily in software would be too slow for such applications. It is
15 important to note that PC and PS/2 computers could also run PHIGS at a satisfactory speed, if a special graphics display board were used.

The artificial reality software that we will use on this project will be able to take the icons to be displayed (with all of the "attribute" information about them--see below) and convert their attributes into a list of polygons that PHIGS should display to the user. This sounds simple, but
20 will actually involve considerable complication. For example, if a group of icons exists at a distance that is too far away to allow the individual icons to be resolved, the software will have to replace the individual icons with a realistic rendering of a "cloud" of icons. In order to allow the effective use of human visual scene analysis, this rendering, and all others, will have to be done with very high quality. Another example is how icons of different size will be displayed at
25 different distances. This will require some clever darkening or shadowing to make size obvious at variable distances. Another possibility is to employ a "fog" that reduces visibility at greater distances. These are all problems that have been solved by artificial reality companies that do

computer-generated graphics for Hollywood films, flight simulators, molecular modeling, high-quality video arcade games, and advertisements.

3.4 Data Object Representation

The DOCUVERSE system will display data objects to the user in the form of three-

- 5 dimensional graphics icons. These icons will have attributes associated with them. An attribute is a user-selected descriptive feature that describes one aspect of the meaning of the data object. Each attribute will have a value between some minimum (say, 0) and some maximum (say, 1). The goal of the display is to show the user all of the attributes of each data object for a large set of data objects in one grand visual display. To do this, we shall exploit almost every human
10 visual (and computer graphics!) skill.

The basic idea is that the icon that represents each individual data object will have many visual degrees of freedom (see FIG. 15).

These degrees of freedom will be "orthogonal", in the sense that they can be set independently of each other. Some examples are: position in space, size, shape, color, distortion,
15 surface texture, and motion. Perhaps, via audio imaging, sound could also be used.

Since the computers used will have graphics capabilities, the presentation of each icon will be visually very sophisticated. Each individually resolvable icon will be portrayed as an easily identified object in its proper position and possessing its assigned attributes. The display will resemble more a photograph of a real scene than an artificial reality display. Except that the
20 objects will be icons. By means of simple mouse commands the user will be able to navigate around in the 3-dimensional cyberspace. A user-selected window will be available to show the entire cyberspace (a cube) and show the user's present position and their current direction of look through the screen of the display. The user will be able to effortlessly change their position to any new (x,y,z) position and rotate their direction of gaze to any new azimuth and elevation (the
25 viewscreen will always remain horizontal, i.e., parallel to the x-y plane, in its boresight rotation). While quick changes in position and angle will be possible, smooth motion will not--as this would require an expensive specialized processor.

Icons that lie at too great a distance from the viewscreen to be resolved will be represented as points of light, or, if they are in groups, as shaded regions--sort of like the milky way. This will allow users to assess the locations of distant data objects. Options such as the introduction of a "fog" that limits distant vision will also be available.

5 At low zoom levels the user will be able to see the entire cubical cyberspace within the viewscreen. At these zoom levels individual objects will typically be invisible and only large groups of documents will be seen. The density of these groups will be represented via 3D shading. By moving the viewscreen around the cyberspace cube the 3-dimensional layout of the entire document set will be visible. This will then suggest strategies for moving into specific
10 regions for a closer look.

Another user option will be to simply double-click on any icon to open that document for inspection. This will cause a new overlay window to appear containing the text of the object and presenting the TIPSTER-like highlighter profiles (see TIPSTER Project documentation for details) for the object (for each of the icon display attributes the user has chosen). By clicking
15 anywhere on any of the highlighter profiles the text block displayed will instantly jump to the designated passage. The ability to roam a document universe at warp speed and then instantly locate and jump to passages of specific content in a document can reasonably be expected to increase analyst productivity many times over. The difference between manual searches through sets of retrieved documents and exploring the same set with the DOCUVERSE system will be
20 like the difference between the Wright Flyer and the Starship Enterprise. We hope to give each analyst the power to continuously monitor all available textual information in their area of responsibility, and still have time to enjoy life.

3.5 Attributes

The attributes that determine the degrees of display freedom of the icons are chosen by
25 the user by making selections via a user interface on the workstation. Preliminary concepts about how these attributes will be selected and used are now discussed. Task 1 of this project will be

devoted to carefully defining the various methods that will be available to the user for selecting and using the attributes.

Coordinate System Selection: Coordinate Words

Spatial coordinates of displayed icons are specified by three context vectors (see FIG. 4).

- 5 Context vectors can be from words ("coordinate words"), groups of words, documents or sets of documents. Coordinate words define a (non-orthogonal) basis space for viewing the projected information content of data objects in 3D. Example coordinate words might be "Terrorist", "Drugs", and "Money Laundering". Positions for display of icons will be computed based on the projection (dot product) of the data objects' context vector with the context vector for the
- 10 coordinate word. It will also be possible to use other data object attributes as "coordinate words". For example, one coordinate axis could be used to display the date of production of each data object. Another could display whether the author was Fidel Castro or not, etc. (See FIG. 16).

Information Content Display Control: Icon Words

- 15 In addition to the coordinate words, the user can specify a set of "icon words". Icon words determine the additional information that will be displayed (i.e., as the shape, color, texture, size, and motion of the icon, see FIG. 17). Icon attributes will be determined by computing the projection of the context vector for the data object attributes with context vector for each icon word. One icon attribute will be associated with each icon word.

Qualification of Display Objects

- 20 Qualification of data objects will be allowed to help reduce visual clutter and information overload. Boolean qualification will be sets of key words to match. Icon attributes will be displayed only for those objects that are a hit. Probably the most common qualification approach will be to simply use a retrieval and routing system to select a body of suitable documents.

Exhaustive enumeration of data objects will be allowed: e.g. Document 4, Document

- 25 1544, Document 3, Database "set.info.lib.text" either as keyboard command or "point and click" with mouse.

Context vectors can be used to perform data object qualification via dot products and thresholding (similar to routing approach). Clustering can be used to associate data objects with similar information content. Only objects within a cluster will be displayed.

Traversing the Space: Changing Point of View

- 5 Selection of coordinate words or object-associated data defines the display basis space. Selection of icon words determines which kinds of information are displayed. Input from a pointing device (mouse, etc.) determines the viewpoint of the user in the 3D coordinate space. Changes in this viewpoint (and in zoom level and angle of view) will be carried out using standard, off-the-shelf computer graphics and artificial reality methods.
- 10 All data objects at the current level of display hierarchy are displayed using 3D, perspective and hidden line removal. User can modify his viewpoint, coordinate or icon specifiers at any time to gain full insight into the information contained in the database (see FIG. 18).

Scope Control: Preventing Information Overload

- 15 Data object qualification provides the first method of overload reduction. Hierarchical organization of the data objects provides the second method of overload reduction. The user selects a data object for examination of subordinate objects using pointing device: "Zoom-in" to lower level of detail. When zooming in on an object, all higher level, non-selected objects are erased from the display. We will consider a mode whereby each document is nominally displayed
- 20 as an icon and where the user will be allowed, if desired, to dive inside that icon and examine the document's chapters, sections and paragraphs as sub-icons. Diving inside a paragraph might cause the text of that paragraph to be automatically displayed in a screen window.

- 25 At the lowest level in the hierarchy the text of the data object is displayed. The user may ascend and descend the hierarchy to any desired level. Higher level objects (such as sets of documents) are restored when ascending the hierarchy.

Automated Determination of High Information Coordinates for use as Attributes

Besides allowing the user to select attributes, the DOCUVERSE system will be capable of automatically selecting coordinates for a document set. These coordinates will be selected on the basis that they have the highest possible information content. For example, the simplest such coordinates would be the principle components of the document set covariance matrix. These are

5 obtained by considering the data object's context vectors as points in feature space. These points form a cloud, with one point for each object. Principal component analysis simply fits the best fitting ellipsoid to this cloud (in a mean squared error sense). The longest perpendicular principal axes of this ellipsoid are then the principal axes of the data set (in many real-world cases only a few of the ellipsoid axes will be very long--the short ones can then be ignored, with little

10 resulting error). The attributes would then be these principal components.

While principal component analysis usually works well as a first-order approximation, it often fails to yield accurate representations. A generalization of the principal component method has been developed by Cottrell, Monro, and Zipser and extended by Hecht-Nielsen. This method uses a neural network to fit a general curvilinear coordinate system inside the data set, in which

15 the coordinate axes remain approximately locally orthogonal, but curve to fit the actual form of the geometrical "cloud" of data in the set. This method is expected to yield attributes that have the highest possible information content in terms of their ability to represent the original document context vectors. HNC is currently engaged in an effort to refine this method. The results of this effort, which is sponsored by the SDIO Office of Innovative Science and

20 Technology and managed by the Mathematical Sciences Division of the Army Research Office, will be used in the project proposed here.

3.6 A Specific Example

Finally, we present a specific fictional example of how some of the capabilities of our proposed DOCUVERSE system might be used. Rather than present an example that explores the

25 more obvious projected capabilities of the DOCUVERSE concept, we discuss a situation that illustrates one of the innovative ways in which the system will probably be used. The objective in

this example is to determine, in near real-time, the nature of the information being transmitted on a high-speed non-encoded (or decoded) data link.

Data: Intercepted high bandwidth (10 MB/sec) data link that contains text information.

Obstacles: There is far too much information for the available analysts to manually read

5 messages within an acceptable amount of time.

Approach:

Step 1: Specify Coordinates

Analyst defines display coordinate space by selecting a set of coordinate words of interest:

10 Drugs

Money Laundering

Terrorists

DOCUVERSE displays representation of information contained in messages as icons in selected coordinate space (see FIG. 19).

15 Step 2: Inspect and Select Data

Analyst determines which messages are of interest in the "terrorist, drugs, money laundering" coordinate system by visual inspection and qualifies (selects) a set of messages with mouse for more detailed inspection (see FIG. 20).

Step 3: Change Point of View

20 Using a mouse, the analyst moves in the selected coordinate space to gain a clearer understanding of information in messages. The display is updated to reflect new point of view (see FIG. 21).

Step 4: Select Icon Words and Attributes

The analyst selects additional data dimensions for investigation by specifying icon words:

25 "Sendero Luminoso"

"Simon Bolivar"

"Cinco de Mayo"

The display is updated to reflect the message information content relative to icon words.

Messages that have high correspondence with these icon words will be easily visible (see FIG. 22).

Step 5: Zoom

5 The analyst will then select the most relevant data object for zoom-in examination of message information. The display is updated to reflect adjusted point of view (see FIG. 23.) 11).

Step 6: Inspect

The text of the selected message is then displayed for manual inspection (see FIG. 24).

3.7 Advantages of the DOCUVERSE System Concept

10 In summary, the DOCUVERSE system will provide new data exploration capabilities for analysts. Some of the advantages of this system are listed below.

High Speed: Allows inspection of the information content of large volumes of isolated text without the need for manual reading of messages.

Cost Effective: Allows fewer analysts to perform more work.

15 Easy to Use: Intuitive user interface exploits natural scene interpretation capabilities of human visual system.

Compatible: DOCUVERSE uses same context vectors and features as HNC's TIPSTER routing/retrieval system.

20 Flexible: No assumptions made about nature of information. Could be applied to interpretation of speech if words are tagged with context vectors. Could even be used for specific-attribute gisting of foreign language message traffic.

Standard Hardware: Standard computer graphics workstations or PC/PS2 computers with a graphics board.

Fast User Training: Intuitive human interface minimizes training requirement.

HNC MatchPlus Functional Specification

HNC MatchPlus Modules

The MatchPlus approach to information retrieval uses high dimensional vectors (280 dimensions) called Context Vectors. Context vectors represent all stems, documents, and queries 5 in the same vector space. This allows direct comparison between any two of these objects by taking a dot product.

System generation begins with the preprocessing of documents to handle format problem, remove stopwords, etc. Next an inverted index is generated for later use in retrieval. The Bootstrap step creates context vectors for all stems, so that similarly used stem have similar 10 vectors. Next, there is the option of generating word senses for commonly used stem, where each word sense receives its own context vector. The final step in system generation is the creation of document context vectors by summing stem context vectors (with appropriate weightings).

For routing queries where there are documents with relevance judgements, a routing data generator extracts data so that neural network algorithm can generate a query. There are two 15 types of routing query, either term weights for a given query or an entire query context vector.

For ad hoc queries, either the user supplies an initial query or the Topic Processor automatically converts a Tipster topic to an initial query.

Once the query has been determined, the Document Retriever fetches as many documents as desired, in order by estimated relevance to the query.

20 As a final step, the Output Combiner may optionally combine several lists of documents retrieved for the same query into one merged list.

The following sections give detailed descriptions of MatchPlus modules in the same order in which they appear in FIGS. 25-28.

1 Document Preprocessor

25 1.1 Functional Description

This module preprocesses raw text, taking stems off of words, taking out stop words, and deleting extraneous characters.

1.2 Data Specifications

1.2.1 Input Specification

Inputs consist of:

The 'raw documents file' in the files of the corpus.

5 The 'stopword file' which lists words to be dropped.

The 'exception file' which lists words that are NOT to be stemmed.

The 'word group file' which gives pairs or groups of words that are to be treated
as a unit during training; an example is 'united states'.

The 'corpus descriptor file' which contains the location of corpus components.

10 1.2.2 Output Specification

Outputs consist of:

The 'preprocessed documents file' which contains stemmed words, word pairs,
etc.

2 Inverted Index Generator

15 2.1 Functional Description

This module makes a pass over the preprocessed documents and creates the inverted
index file and a file of information about each stem in the corpus.

There are two steps in this process:

1. For each stem, the generator creates an inverted index consisting of the stem
and the document(s) that stem occurred in.

20 2. For each stem, the generator creates the 'stem info file' containing information
such as the number of documents the stem occurred in, a pointer to the
context vector for the stem, and the word-sense meaning of the stem.

2.2 Data Specifications

25 2.2.1 Input Specification

Inputs consist of:

The 'preprocessed documents' containing stemmed words and word pairs.

2.2.2 Output Specification

Outputs consist of:

The 'inverted index file' consisting of stems and the documents they appear in.

The 'stem info file' consisting of information about stems such as number of

5 documents the stem occurred in.

3 Stem CV Generator (Bootstrap)

3.1 Functional Description

This module generates an initial random Context Vector for each stem and modifies each stem according to its neighboring stems.

10 Two passes are made through the entire corpus. For each stem in each document the three stems on either side (neighbors) are summed up. After the pass is completed the sum of all the neighbors for each particular stem is factored into that stem's context vector.

3.2 Data Specifications

3.2.1 Input Specification

15 Inputs consist of:

'Preprocessed documents' are needed for getting each stem and allocating a Context Vector (280 floats) and determining the stem targets and neighbors.

The 'system op info file' specifies the document range in which bootstrapping is
20 to occur.

3.2.2 Output Specification

Outputs consist of:

The 'stem info file' with the trained stem CV's.

4 Word Sense Generator

4.1 Functional Description

This module generates word sense information for frequently used stems.

There are three main steps in this process:

1. For each appearance of a frequent stem, a window context vector is computed for surrounding stems. For space considerations, the number of windows for each stem is limited to about 300 randomly selected windows.
2. The windows for each stem are clustered. Initially a variant of k-means clustering is used where the number of clusters is fixed at 7. Each cluster is represented by a centroid CV.
3. For each cluster centroid, the closest 10 stems are found. This information is only useful for communicating with humans and for debugging; it is not used in retrieval.

10 4.2 Data Specifications

4.2.1 Input Specification

Inputs consist of:

The 'stem info file' after bootstrapping. This file contains vectors for all stems.

'Preprocessed documents' are needed for finding window CV's.

The 'system op info file' specifies which part of the text is to be used for collecting windows.

15 4.2.2 Output Specification

Outputs consist of:

The 'word sense file' specifying CV's for every word sense.

20 The 'word sense stem info file', a modified stem info file that contains a pointer into the 'word sense file' for every stem that contains word senses.

5 Document CV Generator

5.1 Functional Description

This module takes the preprocessed documents, stem information, and word sense

25 information and creates document context vectors. The context vector for each stem in a document is weighted according to various formulae, summed together, then normalized to create the document CV.

5.2 Data Specifications

5.2.1 Input Specification

Inputs consist of:

The 'preprocessed documents' containing stemmed words and word pairs.

5 The 'system op info file' which contains information about the range of
documents to process

The 'stem info file' which contains information about how many documents a
stem occurs in, a pointer to its context vector, etc.

5.2.2 Output Specification

10 Outputs consist of:

The 'document context vector file' which contains a vector for each document
derived from the stem context vectors in that document.

6 Routing Data Generator

6.1 Functional Description

15 This module generates the data necessary for stem routing and CV routing.

1. Stem routing data consists of the dot product between each stem in the query and the
judged document. In addition the relevance judgement (a '0' or '1') is entered.
2. CV routing data consists of the document's context vector and the relevance
judgements.

20 6.2 Data Specifications

6.2.1 Input Specification

Inputs consist of:

The Query stems and their CV's.

25 The 'relevance judgements' file with the official document id and the relevance
judgements.

The 'read doc hash table' which is used to convert the official document id from the relevance file to an internal document number. This is necessary to get the document CV.

The 'context cvs doc file' which contains the document CV's.

5 6.2.2 Output Specification

Outputs consist of:

The 'stem routing data file' which is an ascii file with the query stems and the dot products between the query stems and the judged documents.

The 'CV routing data file', which is an ascii file with the products between each query CV vector element and each judged document CV element.

10 7 Query Generator (Learning Algorithms)

15 7.1 Functional Description

This module generates a routing (or ad hoc) query based upon data generated by the Routing Data Generator.

15 The learning algorithm used is the 'pocket algorithm with ratchet', a single-cell neural network algorithm.

20 7.2 Data Specifications

25 7.2.1 Input Specification

For a Stem Routing Query, the inputs are a set of training examples, where each example consists of the relevance judgment for a document, plus the dot product of the document's context vector with the stem context vector for each stem in the original user query.

For a CV Routing Query, the inputs are a set of training examples, where each example consists of the relevance judgment for a document, plus that document's context vector.

25 7.2.2 Output Specification

There are two types of query produced by this module:

A Stem Routing Query consists of a set of weights for each term in the original user query.

A CV Routing Query is a context vector that can be used directly to take dot products with document context vectors in order to determine the order in which documents are retrieved.

5

8 Topic Processor

8.1 Functional Description

This module generates the following from a Tipster Topic:

1. The stems from which to calculate the query CV. The stems may be taken from any of the sections of the topic (e.g. concepts, definitions etc.).
2. The stems that are required to be in a document (i.e. match terms).
3. Overall weights to assign a particular section of the Topic.

8.2 Data Specifications

8.2.1 Input Specification

Inputs consist of:

An ascii Topic file.

User inputs indicating the match criteria which sections of the Topic to use and any weights to apply to a particular section.

The stop word list.

The exception word list.

The word group list.

Word sense stem info file.

Word sense file.

8.2.2 Output Specification

Outputs consist of:

A user query.

10

15

20

25

9 Query Processor

9.1 Functional Description

This module generates a Query CV from a list of words taken from the following:

1. Output from any one of the sections in the Tipster Topics (e.g. concepts section).
2. Entered stems from the GUI

5

9.2 Data Specifications

9.2.1 Input Specification

Inputs consist of:

An ascii file with the query terms.

The stop word list.

10

The exception word list.

The word group list.

Word sense stem info file.

Word sense file.

9.2.2 Output Specification

15 Outputs consist of:

A Query CV (280 real numbers).

10 Document Retriever

10.1 Functional Description

This module takes the output from the Query Processor and produces an ordered list of

20 documents.

First those documents that satisfy the \$Match criteria are identified. (This group may be empty if, for example, no document satisfies the criteria or there are no \$Match criteria.) These documents are then ordered by closeness of their context vectors with the Query Context Vector, and they are retrieved in that order.

25

After the \$Match group is exhausted, remaining documents are ordered by closeness of their context vectors with the Query Context vector, and they are retrieved in that order.

10.2 Data Specifications

10.2.1 Input Specification

The Query Processor determines the \$Match criteria, as well as a Query Context Vector.

5 The 'system op info file' specifies whether all or a subset of documents are eligible for retrieval.

The 'inverted index' allows quick determination of those documents that satisfy the \$Match criteria.

10.2.2 Output Specification

10 Outputs are an ordered list of documents and dot product scores with the Query Context Vector.

11 Output Combiner (Optional)

11.1 Functional Description

This module combines output lists, from MatchPlus or any other source, making use of an estimate of the quality of each source.

15 The algorithm is to give each document a merit score. Letting i run over all input lists, each document receives a score consisting of the sum of the quantity 'goodness of list i'/'position of document in list i'.

Documents are then produced in order of their merit scores.

20 11.2 Data Specifications

11.2.1 Input Specification

Several lists of documents.

An estimated goodness for each list consisting of a non-negative fractional number.

25 11.2.2 Output Specification

An ordered list of documents.

Section II. Image Retrieval Using Context Vectors

In this section II, all reference numbers refer to Figs. 29-43.

- FIG. 29 shows a block diagram of software and hardware components for implementing one embodiment of the present invention 100. Processor 102 is a conventional engineering workstation or other computer processor and may also be coupled to other processors accessible over conventional communications channels or buses (not shown). Processor 102 is conventionally coupled to storage 104 which may be a magnetic disk storage, a CD storage unit, or other conventional computer data storage unit. Storage 104 may also be coupled to other storage units accessible over conventional communications channels or buses (not shown).
- Storage 104 includes image database 106. Image database 106 contains images in electronic or computer-readable form. Images may be discreet images, video, or temporally related images (e.g., motion picture frames).

Processor 102 is also conventionally coupled to memory 108 which is a random access memory (RAM) unit or other conventional computer memory. Items in memory 108 may alternatively be stored in storage 104 and accessed by processor 102 when required. Memory 108 may comprise feature vector generation 110, prototypical feature vector generation 112, iterative training or context vector generation 114, indexing 116, clustering 117, and retrieval 118 program components or modules as discussed below. Input 101 comprises conventional input devices such as a keyboard, mouse, track ball, or touchscreen. A conventional display unit 120 may also be conventionally coupled to processor 102.

The preferred embodiment of the present invention may be implemented on any platform, operating system, and user interface of sufficient power and flexibility, such as: XWindows/MOTIF; Sun/OS SunView; Microsoft Windows, VAX/VMS, and the like, and may be operated in a distributed data / distributed processor environment, if desired. In the preferred embodiment, the invention is capable of handling a very large data corpus, e.g., containing over 10 million images.

This aspect of the present invention is based upon an image representation scheme using context vectors. Context vectors are used to identify images of interest in existing image data bases and/or the routing and dissemination of images of interest in a real-time environment. A context vector is a fixed-length vector having a plurality of component values or weights that are determined based on relationships between meanings of information items or elements. In the present invention, context vectors represent the relative meaning among, or content of, information elements (described below) that are used to define images. The meaning is derived from geometric (e.g., proximity and co-occurrence of information elements in images) and quantitative statistics, data, and analysis (e.g., frequency of occurrence of information elements in the image database). Images are defined using information elements. Context vectors associated with each information element in an image are combined to form a summary vector. Similar images have closely aligned summary vectors, while dissimilar images have roughly orthogonal summary vectors. This representation scheme allows quantification of the similarity of images to be assessed by performing a simple dot product operation on the associated summary vectors, wherein a high dot product indicates a similar meaning.

Note, the absolute orientation of a particular context or summary vector in the vector-space is irrelevant, as long as the relative orientation (with respect to other vectors) is representative of relative proximity of meaning and content. In other words, the problem of finding a set of vectors defining relative meaning has an infinite number of physical solutions in vector-space (absolute orientations) but only one logical solution (relative orientations). The context vector and summary vector generation scheme of the present invention produces a logical solution without regard to the absolute orientation of the vectors.

Context vectors may be associated with words, terms, documents, document portions, queries, quantitative data, or any other type of information item or element. Many techniques for text-based context vector generation are applicable to image processing as well. See, e.g., Section I, Context Generation and Retrieval, above.

For image processing, context vectors are associated with certain image feature vectors. Elements in the feature vector, called features, are derived using affine wavelet transformations on sample points in images in the image database. Features are combined to form a feature vector. Such feature vectors provide an “image vocabulary” analogous to words in text. Due to 5 the large number of possible feature vectors, a set of prototypical feature vectors, or atoms (or image information elements), are derived to form an “atomic vocabulary.” Random high-dimensional context vectors are assigned to each atom. The context vectors are then modified according to the spatial relationship and co-occurrence of the atoms in the images in a procedure called bootstrapping. Compensation for frequency of occurrence of each atom in the database of 10 images can also be made. Summary context vectors for each image are generated by combining the context vectors associated with the atoms that comprise the image.

Context Vector Generation

Referring now to FIG. 30, there is shown a high-level flowchart of the present invention.

15 A corpus of images in image database 106 provides input of images 201 in electronic form to the system. The system is initialized 203 to generate atomic vocabulary 205 and stop list 207. Atomic bootstrapping 209 is then employed to produce context vectors 211 associated with each atom in the vocabulary. Context vectors 211 are used to generate 213 summary context vectors for each image 215. This method is described in more detail below.

20

Atomic Vocabulary Generation

Referring now to FIG. 31, there is shown a flowchart of the process of generating atomic vocabulary 205 according to the present invention. Atomic vocabulary 205 is generated by using a combination of wavelet transformations, self-organization, and conventional statistics as 25 described below. Image features, for example, primitive shapes, textures, or other similar attributes can be identified by wavelet transformations. These features are generally in a specific

spatial orientation to one another in the image. The features, their number, characteristics, and spatial orientations form the context and content information of the image. Thus, features in images are analogous to words in a document. Images can include any number of individual features in any combination. The information content of an image is defined in terms of specific
5 features in a specific spatial orientation and position. When combined in the correct sequence and in the correct spatial positions, the features define the image.

For each image 201 in image database 106, at least one wavelet transformation 301 is made of the image. Wavelets and wavelet transformations 301 are commonly known and are used to provide an efficient approach to localized characterization of information contained
10 within images (see, for example, O. Rioul and M. Vetterli, "Wavelets and Signal Processing", IEEE SP MAGAZINE, October 1991; and, F. Hlawatsch and G.F. Boudreux-Bartels, "Linear and Quadratic Time-Frequency Signal Representations", IEEE SP MAGAZINE, April 1992). In particular, the Gabor wavelet transformation provides a robust representation scheme of features that are sensitive to orientation and localized spatial frequency content. Thus, Gabor wavelets
15 provide a mechanism for characterizing the information content of images in a compact and efficient fashion.

In the present invention, multiple wavelet transformations centered at "sample points" on the image 201 are made 301. To implement this process, Gabor wavelet parameters are selected to capture orientation-sensitive spatial frequencies at a variety of orientation/ frequency
20 combinations over the range of interest. Both sine and cosine wavelets are used to provide phase discrimination. For example, a combination of eight orientations, five spatial frequencies, and both sine and cosine (even and odd order) phases can be used. Thus, a total of $8*5*2$, or 80 wavelets are associated with each sample point on the image 201. The actual number of
25 combinations for a particular application is determined empirically based on the detail and accuracy of image representation required and processing capabilities. Due to the high level of processing power required during wavelet transformation, the preferred embodiment can employ, for example, the Vision Processor (ViP) chip set to perform high-speed image processing

operations. The ViP chip set is available commercially from HNC, Inc., San Diego, CA, as a plug-in circuit board for workstations manufactured by Sun Microsystems, for example.

A simple uniform grid-point approach can be used to determine the "sample points" on the image where the wavelet transformations are centered. Alternatively, more elaborate 5 conventional "fixation control" techniques to determine "points of interest" within images can be used (see, for example, Abbot, Lynn, "A Survey of Selective Fixation Control for Machine Vision" IEEE Control Systems, August 1992). The accuracy of the image representation is linked to both the number of points sampled and the discrimination of the wavelet transformations employed (i.e., the number of coefficients computed at each sample point). In 10 the example above, 80 real values are associated with each sample point. These values comprise a "feature vector" 305. Atoms in an atomic vocabulary 205 are formed from a subset of the feature vectors 305 in the manner described below.

The result of the wavelet transformations 301 and feature vector generation for each sample point on each image can result in an extremely large number of unique feature vectors 15 305. For example, using only a binary representation for the wavelet transformation values could result in 2^{80} (10^{24}) unique atoms. Accordingly, the present invention recognizes that a small fraction of the feature vectors 305 can sufficiently represent the images in the image corpus. Thus, the set of generated features vectors can be reduced to a smaller number of prototypical 20 feature vectors using statistical and neural network techniques.

In one embodiment, a vocabulary size, N, is chosen which defines the problem size 307. A reasonable range for N has been found to be between 10^4 and 10^6 . The exact vocabulary size 25 is empirically determined based on processing power available, the level of detail required for image definition, and the size and diversity of image database 106, for example. Statistical prototype feature vectors are selected using neural network self-organization techniques 309. Basically, the best N feature vectors 305 that represent the entire population of vectors are chosen by using a vector quantization (VQ) training operation. In the present invention, the VQ "learns" 30 statistically representative exemplars given a population of feature vectors 305. Conventional

VQ training techniques 309 such as extensions to clustering techniques (such as the widely used K-means algorithm) or variations on self-organizing approaches (“Self Organization and

- Associate Memory” by Tuevo Kohonen, 2nd Ed, Springer Verlag, 1988 P. 199-209) can be used to transform a potentially infinite number of examples into a finite (but possibly large)
5 atomic vocabulary of statistically representative examples 205. See, e.g., Section I, Context Generation and Retrieval, above, which describes a singular value decomposition (SVD) method for reducing the vector space.

The prototypical feature vectors form atoms or information elements in the atomic vocabulary 205. Each atomic vocabulary item 205 is given a unique atom ID 311. As described
10 below, each atom ID 311 so defined is then associated with a high-dimensional context vector to discern the proximal meaning between such atoms. The resulting VQ network is also saved 313.

In one embodiment, further statistics of the entire set of feature vectors 305 are determined to generate a “stop list.” As in textual applications, images have a significant number of high frequency-of-occurrence atoms that convey little, if any, information. In the textual
15 application analogy, words like “a”, “an”, “but”, and “the”, are put on a “stop list” and not used to define the textual document. Similarly, statistics can be used to define high-occurrence, low informational atoms in images.

Referring now to FIG. 32, there is shown a flowchart of the process of generating stop list 207 according to the present invention. Feature vectors 305 are fed through the previously saved
20 VQ network 313 and are thus quantified or mapped 401 to the nearest-neighbor equivalent atomic vocabulary vector 205. Atomic IDs 311 associated with the mapped atomic vocabulary vector 205 are thus generated 403. Frequency-of-occurrence statistics are computed and accumulated 405 for the atomic IDs 403. Atomic IDs 403 having a frequency exceeding a cutoff
25 407 are placed on “stop list” 207. The cut-off frequency value 407 can be empirically determined based on the image database subject matter and level of image representation detail required, for example. Stop list 207 is used to eliminate low information atoms in when representing images using the atomic vocabulary 205, 311..

Image Characterization and Representation

Any image may be represented in terms of the unique set of information elements in the atomic vocabulary 205, 311. Referring now to FIG. 33, there is shown a flowchart of a method 5 for representing a new image using atomic vocabulary 205. As discussed above in association with feature vector generation (see, 301, 305), feature vectors for a new image 501 are generated by performing a plurality of wavelet transformations at a plurality of sample points 503 to generate a feature vector 507 at each sample point in the new image 501. As in the example above, each feature vector 507 may comprise a vector of 80 real values.

10 The previously saved VQ network 313 is used to map 509 each feature vector 507 to the nearest-neighbor equivalent atomic vocabulary vector 205. This provides a good approximation of the original image 501 and constrains the number of possible combinations of feature vectors 507 used to represent the image 501. Atom IDs associated with each mapped atomic vocabulary vector 205 are thus generated 511. Atoms IDs 511 contained in stop list 207 are removed as they 15 add little information. Thus, an unlabeled, unconstrained image in electronic form 501 is transformed into an image representation 513 comprising atoms 511 from atomic vocabulary 205, 311.

20 The atomic representation 513 of image 501 provides a significant data compression over conventional image representation schemes. In fact, since the atom vocabulary 205 is relatively rich, the atomic representation 513 can often substitute for the original electronic image representation 501 for many purposes. An image can thus be represented by a set of codes 511, as defined by the atomic vocabulary 205, together with an indication of the location to each atom in the image (e.g., locations taken from the sampling grid used).

Context Vector Generation (Bootstrapping)

The context and relationship of atoms 511 in an image and in the entire image database 106 is learned from examples. This learning procedure is referred to as “bootstrapping,” and is basically a type of self organization. The bootstrapping procedure is based on the concept of 5 atom spatial and co-occurrence relationships. That is, atoms that appear in proximity to one another in a significant number of images are conceptually related (just as are words that appear in proximity to one another in a significant number of documents are conceptually related). The closer the atoms are, the stronger the relationship between them. Additionally, the relative importance of the atoms (and their influence on neighbors) is related to both the local frequency 10 (within an image) and the global frequency (within the corpus of images). The specific technique for implementing these relationships is described below.

Referring now to FIG. 34, a method for initializing the system before bootstrapping is shown. According to the present invention, each atom ID 311 (corresponding to atoms in atomic vocabulary 205) is associated 605 with a high-dimensional context vector 601. A high dimensional vector space is advantageous in that 1) initial values for each vector can be 15 generated by using a random Gaussian distribution which results in nearly orthogonal vectors (indicating no initial relationship between context vectors); 2) the information capacity of the vector space is very large when real valued vectors are used; and, 3) vector superposition allows summation of independent vectors into a single vector and later recovery of independent vectors 20 via dot product. Initial values of each element in the context vectors 601 are randomly generated by using a random Gaussian distribution 603. This results in nearly orthogonal vectors indicating no initial relationship between context vectors. Each atom ID 311 is randomly associated 605 with an initialized context vector 603.

Atomic bootstrapping 607 is then performed by a constrained optimization technique 25 where an objective function contains terms for local atom relationships, as well as global terms to insure that the statistics of the entire vector set remains well-behaved. In the preferred embodiment, a commercially available SIMD Neurocomputer Array Processor (SNAP) is used as

a hardware accelerator due to the compute-intensive nature of the learning operation. The SNAP hardware accelerator is available commercially from HNC, Inc., San Diego, CA.

Referring now to FIG. 35, there is shown a flowchart of one embodiment of the bootstrapping procedure 607. Alternative embodiments that are optimized for processing speed 5 may also be employed. In general, the bootstrapping procedure allows all atoms 511 within an image 501 to influence each other to varying degrees depending on proximity, and co- occurrence. The intra-image influence between atoms 511 is distance-weighted according to a Gaussian distribution in that atoms that appear close together in an image are given higher weight. In the preferred embodiment, all atoms 511 influence all other atoms within a single 10 image, although this is not necessary. Alternatively, the influence could be limited to atoms within a certain “radius” of the target atom. This approach is similar to a Parzan window approach that performs non-parametric estimation of probability density based on superposition of Gaussian weighted samples. See, e.g., Parzan, E. “On Estimation of a Probability Density Function,” Annals of Mathematical Statistics 33, 1065-1076, September 1962 and Rosenblatt, M. 15 “Remarks On Some Nonparametric Estimates Of A Density Function,” Annals of Mathematical Statistics, 27, 832-837, 1956. The vector co-occurrence of atoms 511 is learned in proportion to their mutual spatial co-occurrence frequencies across all images 201 in image database 106.

The following procedure is performed during each learning pass. Multiple learning passes through the corpus or database of images are generally required. First 700, the system 20 selects 701 an image 501. Typically, images are selected sequentially, starting with the first image and proceeding through the last image. However, the learning approach makes no assumptions concerning the order of presentation of the images. For the selected image 701, the atoms 511, and associated context vectors (CVs) 605, that comprise the image are retrieved 703. In one embodiment, an image data structure contains a pointer to the linked list of atoms 511 that 25 represent the image 501, and their associated context vectors 605.

An atom 511 of the selected image is targeted 705. Typically, atoms are targeted sequentially, but this is not necessary. The associated context vector for the target atom

(Target_CV(old)) is retrieved 707. A zero-valued temporary context vector for local summation (Temporary_CV) is also created 709.

Each atom 511 in the image 501 that is not the target atom 705 is selected 711 in turn, and the associated context vector of the selected atom is retrieved 713. Alternatively, each atom 5 in a defined “learning region” or “radius” around the target atom 705 can be used (instead of all atoms in the image), for example, all atoms within 2 sample points from the target atom 705. The physical Euclidean distance, D (a scalar), between the target atom 707 sample point and the selected atom 713 sample point is computed 715 (e.g., points and distances defined in sample point grid). This distance is pictorially shown in FIG. 36. The weighted function, G, of D is 10 computed 717 as:

$$G = e^{-D^2/\sigma} \quad (\text{Eq. 14})$$

where σ is the standard deviation of a Gaussian distribution used in the proximity weighting function. The context vector of the selected atom 713 is weighted 719 by G and inversely by the frequency of occurrence of the selected atom 511 in the image 501 (an indicator 15 of importance), f_i . Alternatively f_i is the frequency of occurrence of atom 511 in the entire image database 106. The weighted context vector is added 721 to the temporary context vector 709.

The steps 711 through 721 are repeated 723 for each atom in the image that is not the target atom 705. Once all atoms 711 have been selected 723 (other than target atom 705), the system computes 725 an error vector (Error_CV) for the target context vector 707, according to 20 the equation:

$$\text{Error_CV} = \text{Temporary_CV} - \text{Target_CV}(\text{old}). \quad (\text{Eq. 15})$$

Error_CV may also be frequency weighted in proportion to the frequency of occurrence of the targeted atom in the corpus. Error_CV is squared and added to a variable Mean_SQ_Error 725 to keep track of the total mean squared error metric for the image corpus.

25 The system then updates Target_CV 727 according to the equation:

$$\text{Target_CV}(\text{new}) = \text{Target_CV}(\text{old}) + \alpha * \text{Error_CV} \quad (\text{Eq. 16})$$

where α is a predefined proximity constraint or step size that limits the correcting effect of Error_CV. The variable α is empirically determined based on variances in the image database and a desired learning rate, for example. Additionally, constraints may be added to prevent the target atom from “collapsing in” on the selected atoms. See, e.g., Section I, Context 5 Generation and Retrieval, above. If there are atoms 511 that have not been targeted remaining in the selected image 729, the system returns to step 705. If not, it checks 731 whether there are more images in the corpus. If so, it returns to step 701; if not, the learning pass ends 733. The learning pass may be repeated multiple times. Generally, the process in FIG. 35 is repeated until the rate of change of Mean_SQ_Error 725 is arbitrarily small (i.e., the system is not learning anything new). The cutoff value of Mean_SQ_Error is determined empirically based on variances in the image database and accuracy required. Once the context vectors 727 are iteratively determined, they are saved in storage 104.

The procedure of FIG. 35 is expressed mathematically by the equation:

$$CV_{target}^{new} = CV_{target}^{old} + \sum_{i \neq target}^1 G(\|CV_{target}^{old} - CV_i\|) \bullet CV_i \quad (\text{Eq. 17})$$

where CV_i is the selected context vector and f_i is the frequency of occurrence of atom i in the image (or alternatively the entire image database), and $G()$ is Gaussian weighting $e^{-D^2/\sigma}$.

Summary Vector Generation

20 Context vectors 727 associated with the atoms 511 that comprise an image 501 are combined to form summary vectors for images. Referring now to FIG. 37, there is shown a flowchart of the steps that are performed for each atom 511 in the image 501, in order to generate a summary vector for the image 501. First, 900, for each atom 901 in the image 501, a check is made to determine 902 if atom 901 is on stop list 207. If atom 901 is on stop list 207, the next 25 atom is processed 901; if not, the associated context vector 727 is retrieved 903. Context vector 727 is weighted 905 by the function $1/\ln(N)$ where N is the frequency of occurrence of atom 901

in image database 106, 201 (alternatively, N is the number of images 201 that contain atom 901).

Weighted context vector 905 is accumulated and summed 907. After all atoms 901 for image

5 501 have been processed 909, the resulting summed vector is normalized 911 to produce

normalized summary vector for the image 913. Normalization is performed by dividing each

component in the vector by the absolute magnitude of the context vector. The magnitude of the

context vector is determined by taking the square root of the square of all of the components in

the context vector.

Of course, alternative training techniques may be used. More generalized context vector techniques are disclosed in Section I, Context Generation and Retrieval, above.

10

Summary Vector Storage (Clustering)

Once the summary vectors 913 have been determined, they are stored in storage 104. As

described above in Section I, summary vectors 913 can be stored to further reduce searching time

by creating cluster trees. Referring to Fig. 38, an initial parent node 1001 at the top of a tree

15 indexed as level 0, node 1, contains all of the summary vectors 913 for all images 201 in the

image database 106. A series of child nodes 1005, each branching from the initial parent node

1001, are created at a next level of the cluster tree by employing clustering algorithm 1003. For

each child node 1005, further branching can be made by treating the child as the parent in

clustering algorithm 1003. This process can be repeated until a sufficient level of clustering

20 detail is achieved 1007, the result being centroid consistent clusters 1009.

Fig. 39, further details a flowchart of clustering algorithm 1003. Centroid-consistent clustering algorithm 1003 divides parent summary vectors 1001 among a series of children nodes

1005. A group of clusters is centroid-consistent if every member of every cluster belongs to the

cluster in the group with the closest centroid. A centroid is determined by taking, for each

25 dimension, an average of the component values from all of the summary vectors within the

group. In a preferred embodiment, the convergent k-means clustering algorithm, a well known clustering algorithm, is used. Convergent k-means clustering can be performed as follows:

1. Begin with any initial partition that groups the vectors into k clusters 1101. For example, take the first k summary vectors as single element clusters and set the initial value of the centroid of each cluster to equal its member vector 1103. Assign each of the remaining summary vectors to the cluster having the nearest centroid (computed by dot product with each cluster) and recompute the centroid for the cluster which gains a vector 1105;
2. Take each summary vector in sequence and compute its distance from the centroid of each of the k-clusters 1107 (computed by dot product). If the vector is not currently in the cluster with the closest centroid, move the vector to that cluster and update the centroids of the clusters that gain or lose a summary vector 1109. This step is performed for each summary vector 1111.
3. Repeat step 2 until convergence is achieved, that is until a pass through all of the summary vectors causes no new assignments 1113. Thus, centroid consistent clusters at a new level are created 1115.

Since convergence 1113 may be rather time-consuming to achieve, the clustering algorithm can be simplified by limiting the number of repetitions of step 2 (1107, 1109, 1111). After a given number of repetitions, such as 99, the centroids can be frozen. Then, one or more passes can be made through all of the summary vectors, distributing the vectors to appropriate clusters, but without updating the centroids. The centroids approximated in this manner will no longer be exact centroids, but such approximate centroids may be sufficient for most applications of the present invention. It is not necessary that the centroids be precise according to the present invention; however, it is preferable that the clusters be centroid-consistent with the approximate centroids. “Centroids” as used hereinafter in this application shall mean approximate

centroids—in other words, a centroid sufficient to establish centroid-consistent clusters. Each node is identified by its centroid for use in the searching process.

In forming a next level of clusters 1007, the nodes in the level above become parent nodes to a set of child nodes below. Only the summary vectors assigned to a parent node are 5 used in the clustering algorithm to form the child nodes which branch from that parent. This is repeated across the entire level of parent nodes and on subsequent levels so that fewer and fewer context vectors are assigned to the child nodes on each lower level.

The nodes form a tree pattern in which each node branches from a node in the level above. Thus, each summary vector is assigned to a node on each level of the cluster tree. Each 10 node has a centroid. The bottom-level node assignments for each summary vector and the centroids for each node are stored for use in the search and retrieval algorithms. On the bottom level of the tree, each node points to each summary vector assigned to it. The nodes on the bottom level may be referred to as “buckets.”

Once a cluster tree has been set up 1009, it is a simple matter to add a new record 15 summary vector to the tree. The initial branches of the tree are examined to find the closest centroid. The summary vector is assigned to the node with the closest centroid. Then, the branches from that node are examined for the closest child node centroid, and the process is continued until a bucket is reached. The new record is then assigned to the bucket with the closest centroid of those buckets branching from the node on the previous level to which the 20 summary vector was assigned. The centroids are not changed. This action preserves centroid-consistency of the clusters. If a bucket gets too large, the summary vectors on the bucket can be divided into subclusters on a subsequent level.

Summary Vector Indexing

25 Images can also be indexed to terms. Index terms can be textual words or codes, for example. More than one index term can be associated with an image. For example, an image of

a dog may be indexed to the textual words “dog”, “bark”, and “pet”. Thus, an index to textual terms allows for searching of images using textual retrieval techniques. The underlying principle is that images that have similar content will be associated with similar index terms. Index terms can also be associated with atoms or sets of atoms. The basic indexing algorithm as it relates to

5 text is also described above in Section I.

Referring now to Fig. 40, a high level data flow diagram of image indexing is shown.

First, a set of pre-indexed “training images” is constructed, i.e., associated index terms are determined for each image 1201. The index term/image association is generally derived by humans according to experience. One or more index terms may be associated with each image.

10 An image context system is built according to the invention as described above 1203, i.e., an atomic vocabulary 205 with stop list 207 is constructed, bootstrapping is performed, and summary vectors 913 are generated for each image.

A new image 1204 is characterized 1205 using the existing atomic vocabulary 205 and stop list 207 as described above (i.e., wavelet transformation and mapping to closest atoms in atomic vocabulary), and a summary vector 1206 for the image is generated. Similar images, and their associated index terms, in the existing training corpus are retrieved by calculating a dot product between the new image summary vector 1206 and each summary vectors 913 associated with the images in the training corpus 1207. Optionally, a threshold can be set to retain the best matches, e.g., the top ten similar images and associated index terms 1208. The associated index terms retrieved are weighted by both similarity 1209 and frequency of occurrence 1211 and an ordered list is produced of the most relevant index terms 1213. The new image is then associated with the listed order of index terms 1215. A threshold 1217 can also be used to choose the top N index terms from the weighted list.

Retrieval

Referring now to Fig. 41, there is shown a flowchart of the retrieval method of the present invention. The system allows the user to formulate image queries using several approaches. For example:

5 •*Image Queries*: Whole images can be used as queries. This is, in essence, asking the system to retrieve images that “look like” a particular image. For this query mode, a summary vector for the query image is generated (i.e., using the atomic vocabulary 205, 311) and is used as the query context vector.

10 •*Atomic Queries*: The user can use a mouse or other computer input device to select portions of a displayed image to be used as a query. The system determines the atoms contained in the designated portions of the displayed image and forms a query context vector using the atomic vocabulary 205, 311. The query context vector is computed from a weighted sum of context vectors of the atoms in the portion of the displayed image being used as the query.

15 •*Textual Searches*: As described above textual index terms or codes can be associated with images or atoms. For example, the context vector for the word stem “aircraft” could be aligned with a visual representation of an aircraft. Thus, text-based queries could be used in retrieval of image data.

20 •*Boolean Matches*: The user may specifically identify atomic vocabulary codes to be matched by the search using boolean logic (e.g., and, or, not operators). Boolean logic may also be combined with, for example, the image query technique.

25 Regardless of the approach used, user queries 1301 are converted into a query context vector 1303. For example, images and image portions are represented by the atomic vocabulary vectors 205 and a weighted query vector is formed from the atomic vocabulary vectors (see Fig. 33). Image relevance is assessed by computing the dot product 1305 of each image summary vector 913 with the query context vector 1303, and accumulating the results 1307. The images

are sorted by dot product 1309. Images associated with the summary vectors that produced high dot products are retrieved 1311.

If desired, image summary vectors may be organized in cluster trees for quicker searching. By using the above-described cluster tree storage mechanism (Figs. 38 and 39), the 5 searching task can be greatly accelerated. The query vector is used to identify the centroid that is closest to the query vector (by dot product computation). The search is performed using a depth first tree walk. A branch is followed down the tree, taking the node at each level having the centroid closest to the query vector. The search proceeds down the tree until a bottom level node without children (e.g., a bucket) is reached. Each of the summary vectors in the bucket is 10 compared with the query vector (again by dot product computation) to identify the closest summary vector(s).

The system may also employ relevance feedback, whereby the user specifies which of the retrieved images are most helpful. A new search may then be performed using the summary vector for the specified images. This technique reduces the time required for searches and 15 improves system effectiveness.

Referring to Fig. 42, highlighting of retrieved images may also be provided. Highlighting indicates image portions having high correspondence to the query vector. To perform highlighting, the retrieved image is partitioned into blocks 1401. A context vector is computed for each block 1403, and the dot product of the query context vector 1303 and each block context 20 vector is computed 1405. The block 1401 with the highest dot product has the highest degree of correlation with the query 1303, and is indicated by some visual distinction such as color, inverse video, or some other form of identification 1407. The highlighting concept is pictorially shown in Fig. 43.

Also, as described above, summary vectors of images can be visually displayed for 25 navigation through the corpus of images to find images that are of interest. Essentially, the summary vectors for images and other information items are displayed in a pseudo-orthogonal display having axes corresponding to query terms (the query terms need not actually be

orthogonal to one another). In this application query terms may be atomic vocabulary items, for example. Image summary vectors may also be displayed with visual attributes representing relative similarity of meaning with query terms, e.g., color, shape, orientation, texture.

Practical applications of the present image storage and retrieval method and system are numerous. For example, the system and method can be used in medical imaging to find myopathies, tumors, or bone breaks in medical images and X-rays. Signature recognition of, for example, checks, or other valid signatures can be performed using images of such signatures in image database 106.

As noted above, image database 106 can comprise video or temporally related images.

Thus, video data can be used in the present invention. For example, films can be input into the system frame by frame (or using lower sampling frequencies, e.g., every 5th frame). Thus, films or portions thereof can be compared.

As above, the use of context vectors to represent data from several different data domains in a common universal meaning space is possible (see *Development of Context Vector™* by *Singular Value Decomposition*, above). In the textual application of context vectors, different languages are linked together by word-pairs or “tie” words that have the same meaning in both languages (e.g., water and wasser in English and German, respectively). The context vector of the tie word in English is used as the context vector of the word in the foreign language.

Training of the context vectors in foreign language proceeds normally, except for the fixed context vector associated with the tie words. Context vectors in the foreign language that are close (e.g., dot product) to the tie word context vector will be similar in meaning to the English tie word. Thus, English queries can be used to retrieve foreign works. The joinder of foreign language data domains represents a common “universal meaning space.”

Similarly, a common universal meaning space can be constructed for images (or video) and sound, for example. Certain “ties” can be fixed between images (e.g., a St. Bernard dog) and audio clips (“bark”) (context vectors in the audio domain may be, for example, frequency and/or amplitude measurements). Training of the context vectors in each domain proceeds normally

except for the fixed ties. Thus, a picture of a Chihuahua dog, would be highly linked to the sound clip “bark” by virtue of being close in meaning (i.e., summary vector dot product) to the St. Bernard dog image. Thus, audio and video data domains can be linked in a common or universal meaning space.

5 Similarly, text and images can be linked, for example magazines and pictures in magazines. Here, the text surrounding the image (e.g., in a caption) can be automatically “tied” or associated with the image. The strength of the tie association can also be varied according to spatial proximity of the text to the image, boldness, font, or other factors. See, Section I, Context Generation and Retrieval, above, for an application of context vectors to text. Information
10 elements in the textual application may be, for example, words or word stems (root words).

The above description is included to illustrate the operation to the preferred embodiments and is not meant to limit the scope of the invention. The scope of the invention is limited only by the following claims. From the above discussion, many variations will be apparent to one skilled in the art that are yet encompassed by the spirit and scope of the invention.

We claim:

1. A computer implemented method of generating context vectors representing information elements for retrieval of the information elements or records containing the information elements, the method comprising:

assigning a context vector to each of a plurality of information elements;

initializing the context vectors such that the context vectors are substantially orthogonal to each other in a vector space;

determining proximal co-occurrences of the information elements; and

adjusting the context vectors based on the proximal co-occurrences of the information elements, such that the information elements that frequently proximally co-occur have context vectors with similar orientations in the vector space.

2. The method of claim 1, wherein initializing the context vectors further comprises:

assigning vector components to the vectors using zero-mean, unit-variance Gaussian random number generation.

3. The method of claim 1, wherein a target context vector is a context vector assigned to a target information element, and a neighbor context vector is a context vector assigned to an information element that proximally co-occurs with the target context vector, and wherein adjusting the context vectors comprises:

for each target context vector to be adjusted:

determining an error vector between a target context vector and each neighbor context vector;

updating the target context vector as a function of the error vectors.

4. The method of claim 3, wherein updating the target context vector as a function of the error vectors comprises:

determining a correction vector from the error vectors, where the correction vector is:

$$C_j = \sum_i^{WS} (\|E_{ij}\| - \alpha_{ij}) \hat{E}_{ij}$$

where:

E_{ij} is the error vector $E_{ij} = N_{ij} - T_j$ between the neighbor context vector N_{ij} and a target context vector T_j ;

WS is a window size containing the target context vector and the neighbor context vectors; and

α is a proximity constraint; and

updating the target context vector as

$$T_j^{NEW} = T_j^{OLD} + \frac{\gamma}{F_j} \sum_{i=1}^{F_j} C_j - M$$

where:

γ is a step size;

F_j is the total number of occurrences of information element j ; and

M is a mean context vector for all unique context vectors.

5. The method of claim 1 wherein a target context vector is a context vector assigned to a target information element, and a neighbor context vector is a context vector assigned to an information element that proximally co-occurs with the target context vector, and wherein adjusting the context vectors comprises:

determining a weighted sum vector of neighbor context vectors of a target context vector;

applying the weight sum vector to the target context vector.

6. The method of claim 5, further comprising:

determining the weighted sum vector according to the equation:

$$W_j = \sum_i^{\text{WS}} \frac{G(i)}{D_j} N_{ij}$$

where:

W is the weighted sum vector;

N_{ij} is the neighbor context vector to target context vector T_j;

G(i) is a Gaussian weight for the neighbor context vector i; and

D_j is the number of documents that contain target information element j; and

applying the weighted sum vector to the target context vector according to the equation:

$$T_j^{\text{NEW}} = T_j^{\text{OLD}} + W_j$$

where:

T_j^{NEW} is the updated target context vector; and

T_j^{OLD} is the un-updated target context vector.

7. The method of claim 1 wherein a target context vector is a context vector assigned to a target

information element, and a neighbor context vector is a context vector assigned to an information

element that proximally co-occurs with the target context vector, and wherein adjusting the context

vectors comprises:

determining a weighted sum vector of neighbor context vectors of a target context vector

according to the equation

$$W_j = \sum_i^{\text{WS}} \frac{G(i)}{D_j} N_{ij}$$

where:

W is the weighted sum vector;

G(i) is a Gaussian weight for the neighbor context vector i; and

D_j is the number of documents that contain target information element j;

determining an error vector from the weighted sum vector and the target context vector:

$$E_j = W_j - T_j$$

where

E is the error vector;

T is target context vector;

determining a correction vector C from the error vectors of the neighbor context vectors:

$$C_j = \sum_{i=1}^{F_j} E_i$$

applying the correct vector to the target vector:

$$T_j^{NEW} = T_j^{OLD} + \gamma C_j - M$$

where:

T_j^{NEW} is the updated target context vector; and

T_j^{OLD} is the un-updated target context vector.

γ is a step size; and

M is a mean context vector for all unique context vectors.

8. The method of claim 1, wherein the vector space is defined by a plurality of axes, and the context vector include vector components corresponding to the respective axes, and wherein the axes individually do not have specific semantic associations.
9. The method of claim 1, wherein a target context vector is a context vector assigned to a target information element, and a neighbor context vector is a context vector assigned to an information element that proximally co-occurs with the target context vector, and wherein adjusting the context vectors comprises:
- adjusting the target context vector as a function of the relative importance of at least one of either the target information element or the neighbor context vector with respect to the plurality of information elements.
10. The method of claim 1, wherein a target context vector is a context vector assigned to a target information element, and a neighbor context vector is a context vector assigned to an information element that proximally co-occurs with the target context vector, and wherein adjusting the context vectors comprises:
- adjusting the target context vector as a function of the frequency of occurrence of the target information element; the frequency of occurrence of the neighbor information element; a total number of records containing the target information element, and a total number of record containing the neighbor information element.
11. The method of claim 1, wherein a target context vector is a context vector assigned to a target information element, and a neighbor context vector is a context vector assigned to an information element that proximally co-occurs with the target context vector, and wherein adjusting the context vectors further comprises:
- adjusting the target context vector as a function of:

a distance of the neighbor information element from the target information element, so that neighbor information elements that are closer to the target information element cause the target context vector to be adjusted to be closer to the neighbor context vector; and

a frequency of occurrence of the neighbor information element in records containing the information elements, such that a neighbor information element that is less frequently occurring more strongly causes the target context vector to be adjusted to be closer to that neighbor context vector than a neighbor information element that is more frequently occurring.

12. The method of claim 1, further comprising:

for at least one record comprising a plurality of information elements, determining a summary context vector for the record from the normalized sum of the context vectors of the information elements that comprise the record.

13. The method of claim 12, further comprising:

receiving a query comprising at least one information element;
generating a query context vector from the information element that comprise the query;
retrieving at least one record having a summary context vector with a orientation in the vector space similar to the orientation of the query context vector.

14. The method of claim 13, wherein retrieving at least one record having a summary context vector with a orientation in the vector space similar to the orientation of the query context vector further comprising:

a tree walk of a cluster tree, the cluster tree comprising a hierarchical plurality of nodes, each node having a cluster centroid vector, each cluster centroid vector associated with a cluster of one or more records, and derived from the one or more records contained in the cluster,

the tree walk performed by alternatively selecting a node of the cluster tree that has a centroid cluster vector with a closest orientation in the vector space to the query context vector.

15. The method of claim 1, further comprising:

clustering the context vectors into a plurality of clusters, each cluster having a centroid vector derived from the plurality of context vector contained in the cluster.

16. The method of claim 1, wherein:

information elements having similar meaning have context vectors with similar orientations in the vector space.

17. The method of claim 1, further comprising:

determining a similarity of meaning between a first information element and a second information by performing a vector operation on the context vectors of the first and second information elements.

18. A computer implemented method of generating vectors representing information items for retrieval of the information elements, the method comprising:

selecting a set of R information elements;

determining for the selected set of information elements an RxR mutual co-occurrence matrix based on proximal co-occurrences of the information elements in a plurality of documents; applying Singular Value Decomposition to the mutual co-occurrence matrix to produce a set of first context vectors, the first context vectors having orientations in a D dimensional vector space, where $D \ll R$; and

wherein each first context vector is uniquely associated with one of the selected information elements, and wherein information elements having similar meaning have respective first context vectors with similar orientations in the vector space.

19. The method of claim 13, wherein the mutual co-occurrence matrix comprises for each pair of selected information elements, a normalized measure of the frequency of proximal co-occurrence of the pair of selected information elements.

20. The method of claim 13, wherein the selected set of information elements comprises a first selected set, the method further comprising:

selecting a second set of information elements different from the first selected set of information elements;

associating each of the second set of information elements with a second context vector; and updating the second context vectors at least once using the first context vectors, wherein the first context vectors are fixed during at least one update of the second context vectors.

21. A computer implemented method of retrieving a record from a database containing a plurality of records, each record containing at least one information element having an associated context vector, the method comprising:

for each of a plurality of information elements, storing a context vectors uniquely associated with the information element, the context vectors having the properties that information elements having similar meaning have context vectors with similar orientations in a vector space, and information elements having dissimilar meanings have context vectors with dissimilar orientations in the vector space;

for each of the plurality of records, storing a summary context vector derived from context vectors respectively associated with information elements that comprise the record;

receiving a query;

deriving at least one query information element from the query;

generating a query context vector from the query information element; and

selecting at least one record having a summary context vector with orientation in the vector space that is similar to the orientation of the query context vector.

22. The method of claim 21, wherein selecting at least one record further comprises:
for each of a plurality of records, determining a distance in the vector space between the query
context vector and a summary context vector of a record; and
selecting the record having a least distance between its summary context vector and the query
context vector.
23. A computer implemented of providing a universal meaning space for human understandable
information elements, the method comprising:
selecting a set of first information elements;
creating a first set of context vectors based on proximal co-occurrences of the first information
elements in corpus of records, each first context vector uniquely associated with one of the
first information elements, the context vectors having an orientation in a vector space, such
that first information elements having similar meaning have context vectors with similar
orientations in the vector space;
selecting a set of second information elements, the second information elements different from
the first information elements;
selecting a subset of the first information elements;
for each first information element in the subset, selecting a corresponding second information
element having a human understandable meaning substantially identical to the meaning of
the first information element;
for each of the selected second information elements, associating the second information element
with the context vector of the corresponding first information element;
assigning a context vector to each non-selected second information element; and
adjusting the context vectors of the non-selected second information elements using the context
vectors of the selected second information elements.

24. The method of claim 23, wherein:

the first information elements are words of a first human language, and the second information elements are words of a second, different human language; and

the subset of first information elements and the corresponding second information elements have substantially identical meaning.

25. The method of claim 23, wherein:

the first information elements are symbolic representations of words of a human language encoded in a first data format, and the second information elements are symbolic representations of non-text data encoded in a second data format different from the first format; and

the subset of first information elements and the corresponding second information elements have substantially related meaning even though they have different data formats and different symbolic representations.

26. The method of claim 23, wherein adjusting the context vectors of the non-selected second information elements further comprises:

adjusting the context vectors of the non-selected second information elements, such that non-selected second information elements and selected second information elements having similar meaning have context vectors with similar orientations in the vector space.

27. A computer-implemented process of generating a dictionary of information elements for a database of records, each record including at least one information element, each information element associated with a context vector, each information element having a determinate proximity to other information elements in a record, wherein a neighbor information element is an information element that occurs proximate a target information element in at least one record in the database, the method comprising:

initializing the context vectors associated with information elements in the dictionary, such that initial context vectors are substantially orthogonal to each other in a vector space;

for each information element being a target information element:

selecting neighbor information elements of the target element in at least one record;

modifying the context vector of the target information element using the context vectors of each selected neighbor information elements as a function of the proximity of each neighbor information element to the target information element, and a co-importance of the target information element and the neighbor information element.

28. The method of claim 27, further comprising:

determining the co-importance according to the relative importance of the target information element and the relative importance of the neighbor information element.

29. The method of claim 28, wherein determining the co-importance comprises:

determining a first relative importance of the target information element, inversely according to the frequency of occurrence of the target information element in the records;

determining a second relative importance of the neighbor information element inversely according to the frequency of occurrence of the neighbor information element in records; and

determining the co-importance as a function of the first relative importance and the second relative importance.

30. The method of claim 29, further comprising:

determining the relative importance of an information element by the equation:

$$I_j = B + (1 - B) \left(1 - \frac{\log \left(\frac{1}{ND_j} \right)}{\log \left(\frac{1}{TND} \right)} \right),$$

wherein:

I_j is the relative importance of information element J;

B represents a predefined lower bound;

ND_j represents the number of records containing information element J; and

TND represents the total number of records.

31. The method of claim 27, wherein the co-importance of the target information element and the neighbor information element is determined using the equations:

1) $C_{TN} = I_T I_N$

wherein:

C_{TN} is the co-importance of the target information element and the neighbor information element;

I_T is the relative importance of the target information element; and

I_N is the relative importance of the neighbor information element; and

$$2) I_j = B + (1 - B) \left(1 - \frac{\log \left(\frac{1}{ND_j} \right)}{\log \left(\frac{1}{TND} \right)} \right),$$

wherein:

I_j is the relative importance of an information element J;

B represents a predefined lower bound;

ND_j represents the number of records containing information element J; and

TND represents the total number of records.

32. The method of claim 27, wherein a proximity constraint varies a magnitude of the modification to the context vector of the target information element as a function of both the frequency of occurrence of the target information element and the frequency of the occurrence of each neighbor information element in the records, so that the context vectors of information elements that frequently proximally co-occur do not converge.

33. In a computer system including a storage device containing a plurality of records, each record containing a plurality of information elements, a computer readable medium for configuring and controlling the computer system to generate a plurality of context vectors, the computer readable medium comprising:

an initial context vector generation module, adapted to read and write to the storage device, which initializes to each of a plurality of selected information element an initial context vector, such that the initial context vectors are substantially orthogonal to each other in a vector space, and which writes the initial context vectors to the storage device in association with respective information elements;

a vector training module, adapted to read and write to the storage device, for modifying the context vector of a selected information element, being a target information element, using the context vectors of neighbor information elements that proximally co-occur with the target information element, as a function of the proximity of each neighbor information element to the target information element, and a co-importance of the target information element and the neighbor information element.

34. A method of automatically indexing documents using a defined index of terms, the method comprising:

providing an indexed collection of documents, each document having at least one index term assigned to the document;

providing a plurality terms, including the index terms, each term associated with a context vector, the context vector having the properties that terms having similar meaning have context vectors with similar orientations in a vector space, terms having dissimilar meanings have context vectors with dissimilar orientations in the vector space, and terms which frequently proximally co-occur have context vectors with similar orientations in the vector space; and generating for each indexed document a context vector from the context vectors of selected terms that comprise the document;

receiving a new document to be indexed;

generating a new context vector of the new document, the new context vector generated from the context vectors of selected terms that comprise the new document;

selecting at least one indexed document having a context vector similar to the new context vector;

assigning to the new document at least one index term assigned to a selected indexed document.

35. The method of claim 34, wherein assigning to the new document at least one index term further comprises:

for each selected indexed document, assigning a weight to each index term assigned to the indexed document, the weight proportional to the similarity between the new context vector and the context vector of the indexed document, such that the weight is higher where the context vectors are more similar;

for each index term, generating an index term score as a function of a number of occurrences of the index term in each selected indexed document, and the weight of the index term with respect to each selected indexed document, such that the index term score is higher as the number of occurrences of an index term increases; and

assigning to the new document at least one of the index terms with a high index term score.

CONTEXT VECTOR GENERATION AND RETRIEVAL

ABSTRACT OF THE DISCLOSURE

A system and method for generating context vectors for use in storage and retrieval of documents and other information items. Context vectors represent conceptual relationships among information items by quantitative means. A neural network operates on a training corpus of records to develop relationship-based context vectors based on word proximity and co-importance using a technique of "windowed co-occurrence". Relationships among context vectors are deterministic, so that a context vector set has one logical solution, although it may have a plurality of physical solutions. No human knowledge, thesaurus, synonym list, knowledge base, or conceptual hierarchy, is required. Summary vectors of records may be clustered to reduce searching time, by forming a tree of clustered nodes. Once the context vectors are determined, records may be retrieved using a query interface that allows a user to specify content terms, Boolean terms, and/or document feedback. The present invention further facilitates visualization of textual information by translating context vectors into visual and graphical representations. Thus, a user can explore visual representations of meaning, and can apply human visual pattern recognition skills to document searches.





FIG. 1A

FIG. 1B






FIGURE 2A




FIGURE 2B




FIGURE 2C




FIGURE 2D





FIGURE 2E



FIGURE 2F

FIG. 3






FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9A




FIG. 9B




FIG. 9C

000000000000000000000000





FIG. 10

FIG. 11






FIG. 12




FIG. 13




FIGURE 14




FIG. 15




FIG. 16

Association
Icon Word

	Strong	Weak	
Sendero Luminoso	○ □	○ □	Sendero Luminoso attribute=Size
Simon Bolivar	○ □	○ □	Simon Bolivar attribute=Distortion

FIG. 17




FIG. 18




FIG. 19




FIG. 20




FIG. 21




FIG. 22




FIG. 23




FIG. 24




FIGURE 25




FIGURE 26




FIGURE 27

Routing Query Generation




FIGURE 28

Retrieval




FIG. 29




FIG. 30




FIG. 31




FIG. 32




FIG. 33




FIG. 34




FIG. 35




FIG. 36




FIG. 37




FIG. 38




FIG. 39




FIG. 40




FIG. 41




FIG. 42




FIG. 43

**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION**

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name:

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

Improved System and Method of Context Vector Generation and Retrieval

the specification of which (check one) is attached hereto.

was filed on _____
Application Serial No. _____
and was amended on _____
(if applicable)

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, §1.56(a).

I hereby claim foreign priority benefits under Title 35, United States Code §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)	Priority Claimed
N/A (Number) _____ (Country) _____	<input type="checkbox"/> Yes <input type="checkbox"/> No <input type="checkbox"/> Yes <input type="checkbox"/> No
(Day/Month/Year Filed) _____	
(Number) _____ (Country) _____	
(Day/Month/Year Filed) _____	

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, §1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

(Application Serial No.) (Filing Date) (Status) (patented, pending, abandoned)

(Application Serial No.) (Filing Date) (Status) (patented, pending, abandoned)

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

Elizabeth F. Enayati, Reg. No. 34,327
Dennis S. Fernandez, Reg. No. 34,160
David L. Hayes, Reg. No. 34,838
Kenneth M. Kaslow, Reg. No. 32,246
Stuart P. Meyer, Reg. No. 33,426

Edward J. Radio, Reg. No. 26,793
A. C. Smith, Reg. No. 20,355
Amir H. Raubvogel, Reg. No. 37,070
Phong K. Truong, Reg. No. P-37,499
Leo Novakoski, Reg. No. P-37,198

SEND CORRESPONDENCE TO:

Amir H. Raubvogel, Esq.
FENWICK & WEST
Two Palo Alto Square, Suite 500
Palo Alto, CA 94306

DIRECT TELEPHONE CALLS TO:
(name and telephone number)

Amir H. Raubvogel, Esq.
(415) 858-7153

Case Docket No. _____

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full name of sole or first inventor William Robert Caid

Inventor's signature WR Caid Dated: 9-17-93

Residence 5438 Jamestown Road Citizenship U.S.A.

Post Office Address (same as residence)

Full name of 2nd joint inventor, if any Joel Lawrence Carleton

2nd Inventor's signature JL Carleton Dated: 9/17/93

Residence 12628 Torrey Bluff Drive, #328 Citizenship U.S.A.

Post Office Address (same as residence)

Full name of 3rd joint inventor, if any Pu Qing

3rd Inventor's signature 1. P. Q. I. Dated: 9/17/93

Residence 2347 Terraza Ribera Citizenship The People's Republic

Post Office Address La Costa, CA 92130 Citizenship of China

Post Office Address (same as residence)

Full name of 4th joint inventor, if any David John Sudbeck

4th Inventor's signature David Sudbeck Dated: 9/17/93
Residence 2542 Mitten Lane Citizenship U.S.A.

Post Office Address (same as residence)

Full name of 5th joint inventor, if any

5th Inventor's signature Dated:

Residence Citizenship

Post Office Address

IN THE
UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICANTS: William R. Caid, Pu Qing, Joel L. Carleton, and David J. Sudbeck

APPLICATION NO.: Please assign

FILING DATE: September 27, 2000

TITLE: Context Vector Generation & Retrieval

EXAMINER: Unknown

GROUP ART UNIT: Unknown

ATTY. DKT. NO.: 5382 US

COMMUNICATION TO EXAMINER

Dear Sir:

Please note that this application is a continuation of U.S. Serial No. 08/931,927, filed September 17, 1997, which was a continuation of U.S. Serial No. 08/322,313 filed October 13, 1994, which was a continuation-in-part of U.S. application Serial No. 08/124,098 filed by Caid et al., on September 20, 1993.

The specification as submitted herewith contains the entirety of the disclosure of both the 08/322,313 (grandparent) application and the 08/124,098 (great-grandparent) application. Support for disclosure of 08/124,098 is based on the fact that the 08/322,313 application as filed incorporated the 08/124,098 application by reference.


The incorporation of the 08/124,098 disclosure into the body of the instant specification has resulted in the renumbering of the Figures of the grandparent application as Figs. 29-43, and renumber of Equations 1-4 of the grandparent application as Equations 14-17. In addition, minor grammatical changes have been made to provide for ease of reading. Accordingly, no new matter is added by incorporation of the great grandparent application's disclosure.

Finally, as the subject matter of the instant specification now has a priority date of September 20, 1993, the Examiner is respectfully requested to conduct his search with respect to this priority date.

If the Examiner has any questions about the foregoing, he is invited to contact the undersigned attorney by telephone.

Respectfully submitted,
WILLIAM CAID & PU QING

By:


Robert R. Sachs

Registration No. 42,120
Fenwick & West LLP
Two Palo Alto Square
Palo Alto, CA 94306
(415) 875-2410 (Phone)
(415) 281-1350 (Fax)